


AFRL-IF-RS-TR-1998-6
Final Technical Report
March 1998



COMPOSING MEGAMODULES AND MEGAPROGRAMS

SRI International

Sponsored by
Advanced Research Projects Agency
ARPA Order No. A712

19980505 129

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Advanced Research Projects Agency or the U.S. Government.

DTIC QUALITY INSPECTED 4

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-IF-RS-TR-1998-6 has been reviewed and is approved for publication.

APPROVED:

Joseph A Carozzoni
JOSEPH A. CAROZZONI
Project Engineer

FOR THE DIRECTOR:

Northrup Fowler
NORTHRUP FOWLER III
Technical Advisor
Information Technology Division

If your address has changed or if you wish to be removed from the Air Force Research Laboratory Rome Research Site mailing list, or if the addressee is no longer employed by your organization, please notify AFRL/IFTB, 525 Brooks Road, Rome, NY 13441-4505. This will assist us in maintaining a current mailing list.

Do not return copies of this report unless contractual obligations or notices on a specific document require that it be returned.

COMPOSING MEGAMODULES AND MEGAPROGRAMS

Robert Riemenschneider

Contractor: SRI International

Contract Number: F30602-93-C-0245

Effective Date of Contract: 16 July 1993

Contract Expiration Date: 2 January 1997

Program Code Number: 5E30

Short Title of Work: Composing Megamodules and Megaprograms

Period of Work Covered: Jul 93 - Jan 97

Principal Investigator: Robert Riemenschneider

Phone: (415) 326-6200

AFRL Project Engineer: Joseph A. Carozzoni

Phone: (315) 330-7796

Approved for public release; distribution unlimited.

This research was supported by the Advanced Research Projects Agency of the Department of Defense and was monitored by Joseph A. Carozzoni, AFRL/IFTB, 525 Brooks Road, Rome, NY.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188
<p>Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.</p>			
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED	Final Jul 93 - Jan 97
4. TITLE AND SUBTITLE		5. FUNDING NUMBERS	
COMPOSING MEGAMODULES AND MEGAPROGRAMS		C - F30602-93-C-0245 PE - 61101E PR - A712 TA - 00 WU - 01	
6. AUTHOR(S)		8. PERFORMING ORGANIZATION REPORT NUMBER	
Robert Riemenschneider		N/A	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)		10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
SRI International 333 Ravenswood Avenue Menlo Park CA 94025		AFRL-IF-RS-TR-1998-6	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		11. SUPPLEMENTARY NOTES	
Advanced Research Projects Agency 3701 North Fairfax Drive Arlington VA 22203-1714		Air Force Research Laboratory/IFTB 525 Brooks Road Rome NY 13441-4505	
12a. DISTRIBUTION AVAILABILITY STATEMENT		12b. DISTRIBUTION CODE	
Approved for public release; distribution unlimited.			
13. ABSTRACT (Maximum 200 words)			
<p>This report describes the research performed to design and develop a formal definition of software architectures in support of system composition. Documented are new techniques and tools to automate the composition of large, parallel, and/or distributed software systems from existing, traditionally constructed modules. Also supported is the execution of assembled systems on a variety of hardware architectures. The techniques and automated tools were successfully used to construct architecture descriptions consisting of over one million lines of source code.</p>			
14. SUBJECT TERMS			15. NUMBER OF PAGES
Software, Knowledge-Based Systems			102
			16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

Contents

1	Introduction	2
2	Formal Architecture Definition	3
3	Results of this Effort	5
3.1	SRI	5
3.2	Stanford University	6
4	Conclusions	7
A	SRI's X/Open DTP Specification	8
B	SRI Publication: Correct Architecture Refinement	62
C	SRI Publication: Correctness and Composition of Software Architectures	80

1 Introduction

This report describes the work performed by SRI International and its subcontractor Stanford University on contract F30602-93-C-0245, dealing with formal definition of software architectures to support system composition. Our basic approach to architecture definition is presented in Section 2, while Section 3 outlines the work performed on this particular project. The details of our efforts can be found in the appendices.




Figure 1: The SRI-Stanford Approach to Architecture Definition

2 Formal Architecture Definition

The Software Architectures team at the SRI Computer Science Laboratory (CSL) and the Program Analysis and Verification Group (PAVG) at Stanford University are engaged in a joint effort to develop concepts and tools for formally defining software architecture hierarchies. Figure 1 illustrates our approach.

The box in the upper left corner of the figure illustrates the structure of a typical SADL architecture hierarchy. The most abstract specification, $S_{\langle 0 \rangle}$, is the root of a tree in which each node is a SADL architecture specification and each arrow is a SADL mapping. An architecture hierarchy need not be a tree. Any partial order is perfectly acceptable. But developing an architecture hierarchy by top-down refinement will produce a tree. The three successors of $S_{\langle 0 \rangle}$ — i.e., $S_{\langle 0 \rangle}$, $S_{\langle 1 \rangle}$, and $S_{\langle 2 \rangle}$ — represent three alternative ways of making the abstract $S_{\langle 0 \rangle}$ architecture somewhat more concrete. Focusing on the leftmost branch of the tree, architecture $S_{\langle 00 \rangle}$ is a further refinement of architecture $S_{\langle 0 \rangle}$, architecture $S_{\langle 000 \rangle}$ is a further refinement of architecture $S_{\langle 00 \rangle}$, and so on down to $S_{\langle 000 \rangle\alpha}$, an implementation-level architecture that is a refinement of all its ancestors. Generally, in this tree-shaped hierarchy, specifications are indexed so that S_{σ} is an ancestor of S_{τ} if and only if σ is an initial subsequence of τ (i.e., if and only if, for some finite sequence ρ , $\tau = \sigma \cdot \rho$).

The box in the upper right corner of the figure contains a pair of executable Rapide architecture protocol simulations. Each Rapide architecture corresponds

to one of SADL architectures, as indicated by the heavy doubleheaded arrows (and matching indices). This correspondence is not formally specified. Each architectural specification contains information the other does not. The Rapide specification contains behavioral information required for simulation that is typically irrelevant to the SADL structural specification, and which is therefore omitted. The SADL specification encodes details about the logical strength of the architectural styles being employed, details that are crucial to the analysis of refinement correctness, and that are not expressible in Rapide. Someone familiar with both languages can easily judge whether a SADL architecture and a Rapide architecture “correspond”, in other words, whether they consistently describe a system, at the same level of abstraction, but from differing perspectives. For these reasons, formalizing the correspondence — as opposed to relying on convention, such as using the same name for corresponding components — would be of little utility.

The two Rapide architectures in the figure are linked by a Rapide event mapping. Any number of SADL architectures can have corresponding Rapide architectures. This event mapping is partially determined by composing the SADL mappings that link the corresponding SADL architectures. In the figure, the architectural protocols are simulated at both an abstract level in $X_{(2)}$ and at a quite concrete (implementation) level in $X_{(201),\delta}$.

The lower box shows implementations of some of the most concrete SADL architectures, linked to their specifications by a mapping expressed in a programming language-specific extension of SADL’s mapping language. (An extension for Java is under development.) In this example, only two of the five implementation-level SADL architectures have been instantiated as code. The dashed arrow from $X_{(201),\epsilon}$ to $I_{(201),\epsilon}$ indicates a nonformalized mapping of the Rapide simulation of the architectural protocols to an implementation of those protocols in the instantiation. The feasibility of replacing this dashed arrow by automatic code generation — based on an implicit formal mapping — is under investigation.

The approach to formal definition of architectures described above provided the foundation for the research performed for this project. More detail, including motivations for creating a hierarchy and the advantages of doing so, can be found in the appendices (Appendix B, in particular).

3 Results of this Effort

3.1 SRI

Prior to this project, SRI CSL had developed a formal architecture definition system, called *PegaSys*, for a commercial client. *PegaSys* addresses a very special case of the general problem addressed by SADL. In *PegaSys*, only two styles of architectural specification were supported. An architecture could be specified at an abstract level using a dataflow style, or at a concrete level using reading and writing of arrays of variables and control signals to implement dataflow. These two styles were already being used for informal architectural specification by the customer. *PegaSys* hierarchies thus had a very simple, restricted structure: refinements either replace a component by a collection of connected components (“bubble decomposition”) or implement dataflow. *PegaSys* tools checked

- the syntactic correctness of specifications,
- whether type constraints on connections were satisfied, and
- whether refinements could be verified by creating a combination of some hardwired refinement patterns that matched the refinement step.

This system proved useful in practice. Several bugs were found in the architectural descriptions of large (100,000 to 1,000,000 lines of source code) control systems by formalizing those descriptions in the *PegaSys* language and checking them with the *PegaSys* tools.

The main emphasis in the present project was on generalizing *PegaSys* to deal with other domains — additional architectural styles, more complicated hierarchies, and so on — and replacing the ad hoc, informal notion of hierarchy correctness employed in *PegaSys* by a more precise criterion. *PegaSys* specifications can be converted to SADL specifications with relatively little change, but SADL is a far richer language. In addition to particular architectures, SADL can be used to define *constraints*, *generic architectures*, *styles*, *mappings* between architectures and between styles, and *refinement patterns*. See the SADL manual, available on the web at

[<http://www.cs1.sri.com/sadl/sadl-intro.ps.gz>](http://www.cs1.sri.com/sadl/sadl-intro.ps.gz)

for details.

The SADL extensions were driven by an analysis of examples. Both simple particular architectures, such as the compiler architecture used in the paper included as Appendix B of this report, and complex generic architectures, such as X/Open’s Distributed Transaction Processing (DTP) standard architecture described in Figure 2, were formalized in SADL. The result of formalizing X/Open DTP has been included as Appendix A.

Once the language design stabilized, tool development began. The SADL 1.0 software distribution, available at

[<http://www.cs1.sri.com/sadl/sadl-distribution.tar.gz>](http://www.cs1.sri.com/sadl/sadl-distribution.tar.gz)




Figure 2: The X/Open DTP Architecture Hierarchy

contains a parser, printer, and mapping checker for the language. See

[<http://www.csl.sri.com/sadl/README.html>](http://www.csl.sri.com/sadl/README.html)

to get an idea of the toolset's present capabilities.

The result of our attempt to define a formal correctness criterion for architecture structure hierarchies can be found in Appendix B. (Appendix C shows how an external semantics can be provided for connector types, which can be useful both for explanation and for showing the consistency of the SADL constraints that internally define a connector type.)

3.2 Stanford University

Prior to this effort, Stanford PAVG developed the Rapide language as a general simulation tool. On this project, PAVG researchers showed how Rapide can be used for architectural definition, by formalizing complex architectures, and extended the capabilities of the toolset.


The Rapide toolset can be found at

[<http://anna.stanford.edu/rapide/tools-release.html>](http://anna.stanford.edu/rapide/tools-release.html)

4 Conclusions


Our principal objective in this contract was used to demonstrate the utility of the basic approach to architecture definition described in Section 2. We believe that this objective was satisfied by fully formalizing a complex architectural hierarchy involving features — such as a variable number of components — that other architecture definition languages cannot handle in any straightforward fashion, by formalizing the notion of hierarchy correctness so that the precise benefits of correctness are clear, and by developing our toolsets to the point where they can be used by others interested in experimenting with formal architecture definition.

A SRI's X/Open DTP Specification



Specification Hierarchy Development Stages

9



(2)

```
%%% Starting point is a very abstract view that treats the collection of
%%% resource managers as a single component

x_open_abstract_top: ARCHITECTURE [ -> ]

IMPORTING ALL FROM Dataflow_Relations_style

BEGIN

CONFIGURATION

%% Note that the following aren't really component declarations, since
%% there is no signature declared. We're at a more abstract level, where
%% all we're saying is that there are components called the_ap, et al.,
%% of some type such that . . .

ap: TYPE <= Function
rms: TYPE <= Function
tm: TYPE <= Function

the_ap: ap
the_rms: rms
the_tm: tm

ar: CONSTRAINT = Dataflow(the_ap, the_rms)
tx: CONSTRAINT = Dataflow(the_ap, the_tm)
xa: CONSTRAINT = Dataflow(the_tm, the_rms)

END x_open_abstract_top
```

```

%% First step is to go to a style that makes the dataflow connections
%% explicit

x_open_abstract_df: ARCHITECTURE [ -> ]

IMPORTING ALL FROM Dataflow_style

BEGIN

ar_requests: TYPE
ar_resources: TYPE
tx_commands, tx_responses: TYPE
xa_commands, xa_responses: TYPE

COMPONENTS

ap:  TYPE <= Function [ap_in1: ar_resources, ap_in2: tx_responses
                      -> ap_out1: ar_requests, ap_out2: tx_commands]

rms: TYPE <= Function [rm_in1: ar_requests, rm_in2: xa_commands
                      -> rm_out1: ar_resources, rm_out2: xa_responses]

tm:  TYPE <= Function [tm_in1: tx_commands, tm_in2: xa_responses
                      -> tm_out1: tx_responses, tm_out2: xa_commands]

the_ap: ap
the_rms: rms
the_tm: tm

%% No named connectors, due to parameterization, hence no CONNECTORS section

CONFIGURATION

ar_1: CONNECTION =
      (EXISTS c: Channel<ar_requests>
       Connects(c, the_ap.ap_out1, the_rms.rm_in1)
ar_2: CONNECTION =
      (EXISTS c: Channel<r_resources>
       Connects(c, the_rm.rm_out1, the_ap.ap_in1)

tx_1: CONNECTION =
      (EXISTS c: Channel<tx_commands>
       Connects(c, the_ap.ap_out2, the_tm.tm_in1)
tx_2: CONNECTION =
      (EXISTS c: Channel<tx_responses>
       Connects(c, the_tm.tm_out1, the_ap.ap_in2)

xa_1: CONNECTION =
      (EXISTS c: Channel<xa_commands>
       Connects(c, the_tm.tm_out2, the_rms.rm_in2)
xa_2: CONNECTION =
      (EXISTS c: Channel<xa_responses>
       Connects(c, the_rms.rm_out2, the_tm.tm_in2)

END x_open_abstract_df

```

```

%%% Replace the aggregate resource managers component with a
%%% ARCHITECTURE that contains the individual resource managers. Although this
%%% is complicated, it seems to be just two xformations, one applied twice.
%%% First, the ports and channels are split. Second, the Function is
%%% replaced by a ARCHITECTURE. I suppose an "empty" ARCHITECTURE could be introduced
%%% and then refined by adding the processes -- which has to be done all
%%% at once when there is no particular number of them --, which is what
%%% the rule in the paper suggests, but that just complicates analysis of
%%% ARCHITECTURE interface constraints.

```

```

%%% After the two channel splitting, we have
x_open_intermediate_1_df: ARCHITECTURE [ -> ]

```

```
IMPORTING ALL FROM Dataflow_style
```

```
BEGIN
```

```

n: NAT      % Number of resource managers, a parameter in the specification

ar_requests, ar_resources: TYPE
tx_commands, tx_responses: TYPE
xa_commands, xa_responses: TYPE

%% q_type(i) will be the subtype of ar_requests accepted by the i-th resource
%% manager, and similarly for r_type(i) and ar_resources.
q_type: {i: NAT | i < n} --> {t: TYPE | t < ar_requests}
r_type: {i: NAT | i < n} --> {t: TYPE | t < ar_resources}

```

```
COMPONENTS
```

```

ap: TYPE <= Function [ << ap_in1(i): r_type(i) |(i: NAT) i < n >>,
                      ap_in2: tx_responses
                      -> << ap_out1(i): q_type(i) |(i: NAT) i < n >>,
                         ap_out2: tx_commands]

rms: TYPE <= Function [<< rm_in1(i): q_type(i) |(i: NAT) i < n >>,
                        << rm_in2(i): xa_commands | (i: NAT) i < n >>
                        -> << rm_out1(i): r_type(i) |(i: NAT) i < n >>,
                           << rm_out2(i): xa_responses | (i: NAT) i < n >>]

tm: TYPE <= Function [tm_in1: tx_commands,
                      << tm_in2(i): xa_responses | (i: NAT) i < n >>
                      -> tm_out1: tx_responses,
                         << tm_out2(i): xa_commands | (i: NAT) i < n >>]

the_ap: ap
the_rms: rms
the_tm: tm

```

```
CONFIGURATION
```

```

ar_1: CONNECTION =
  (FORALL i: NAT | i < n)
    (EXISTS c: Channel<q_type(i)>)
      Connects(c, the_ap.ap_out1(i), the_rms.rm_in1(i))
ar_2: CONNECTION =
  (FORALL i: NAT | i < n)
    (EXISTS c: Channel<r_type(i)>)
      Connects(c, the_rms.rm_out1(i), the_ap.ap_in1(i))

tx_1: CONNECTION =
  (EXISTS c: Channel<tx_commands>)
    Connects(c, the_ap.ap_out2, the_tm.tm_in1)
tx_2: CONNECTION =

```

```
(EXISTS c: Channel<tx_responses>)
    Connects(c, the_tm.tm_out1, the_ap.ap_in2)

xa_1: CONNECTION =
    (FORALL i: NAT | i < n)
        (EXISTS c: Channel<xa_commands>)
            Connects(c, the_tm.tm_out2(i), the_rms.rm_in2(i))
xa_2: CONNECTION =
    (FORALL i: NAT | i < n)
        (EXISTS c: Channel<xa_responses>)
            Connects(c, the_rms.rm_out2(i), the_tm.tm_in2(i))

END x_open_intermediate_1_df
```

```
%%% Refining the Function the_rms into a ARCHITECTURE containing many rm's yields
x_open_intermediate_2_df: ARCHITECTURE [ -> ]
```

```
IMPORTING ALL FROM Dataflow_style
```

```
BEGIN
```

```
n: NAT % Number of resource managers, a parameter in the specification
```

```
ar_requests, ar_resources: TYPE
tx_commands, tx_responses: TYPE
xa_commands, xa_responses: TYPE
```

```
%% That the q_type is a partition of ar_requests is guaranteed by the
%% general constraints on ARCHITECTURE interfaces and the "completeness
%% assumption". (We say nothing about the ports of the resource
%% managers -- in particular, we mention no connections withing the
%% ARCHITECTURE -- so all are externally visible.) Ditto for r_type and
%% ar_resources.
```

```
q_type: {i: NAT | i < n} --> {t: TYPE | t < ar_requests}
r_type: {i: NAT | i < n} --> {t: TYPE | t < ar_resources}
```

```
COMPONENTS
```

```
ap: TYPE <= Function [<< ap_in1(i): r_type(i) |(i: NAT) i < n >>,
                      ap_in2: tx_responses
                     -> << ap_out1(i): q_type(i) |(i: NAT) i < n >>,
                      ap_out2: tx_commands]
```

```
rm: TYPE <= { p: Function[rm_in1: qt, rm_in2: xa_commands
                           -> rm_out1: rt, rm_out2: xa_responses]
                  | qt < ar_requests AND rt < ar_resources }
```

```
rms: TYPE <= ARCHITECTURE [<< rm_in1(i): q_type(i) |(i: NAT) i < n >>,
                            << rm_in2(i): xa_commands | (i: NAT) i < n >>
                           -> << rm_out1(i): r_type(i) |(i: NAT) i < n >>,
                            << rm_out2(i): xa_responses | (i: NAT) i < n >>]
```

```
tm: TYPE <= Function [tm_in1: tx_commands,
                      << tm_in2(i): xa_responses | (i: NAT) i < n >>
                     -> tm_out1: tx_responses,
                      << tm_out2(i): xa_commands | (i: NAT) i < n >>]
```

```
the_ap: ap
the_rms: rms
the_tm: tm
```

```
CONFIGURATION
```

```
rms_contents: CONSTRAINT =
  (FORALL y: COMPONENT)
    [y PROPERLY_CONTAINED_IN the_rms => (EXISTS z: rm) y CONTAINED_IN z]
rm_location: CONSTRAINT =
  (FORALL y: rm) [y CONTAINED_IN the_rms]
```

```
ar_1: CONNECTION =
  (FORALL i: NAT | i < n)
    (EXISTS c: Channel<q_type(i)>
      Connects(c, the_ap.ap_out1(i), the_rms.rm_in1(i))
```

```
ar_2: CONNECTION =
  (FORALL i: NAT | i < n)
    (EXISTS c: Channel<r_type(i)>
      Connects(c, the_rms.rm_out1(i), the_ap.ap_in1(i))
```

```
tx_1: CONNECTION =
  (EXISTS c: Channel<tx_commands>)
    Connects(c, the_ap.ap_out2, the_tm.tm_in1)
tx_2: CONNECTION =
  (EXISTS c: Channel<tx_responses>)
    Connects(c, the_tm.tm_out1, the_ap.ap_in2)

xa_1: CONNECTION =
  (FORALL i: NAT | i < n)
    (EXISTS c: Channel<xa_commands>)
      Connects(c, the_tm.tm_out2(i), the_rms.rm_in2(i))
xa_2: CONNECTION =
  (FORALL i: NAT | i < n)
    (EXISTS c: Channel<xa_responses>)
      Connects(c, the_rms.rm_out2(i), the_tm.tm_in2(i))

END x_open_intermediate_2_df
```

```

%%% Refine the TX interface by splitting the tx_command,
%%% tx_response, xa_command, and xa_response channels to set up use
%%% of the "actual" commands .

x_open_concrete_df: ARCHITECTURE [ -> ]

IMPORTING ALL FROM Dataflow_style

BEGIN

n: NAT      % Number of resource managers, a parameter in the specification
ar_requests, ar_resources: TYPE

%% Arguably, the whole TX interface is control flow, but we'll treat
%% the integers that get returned as data to stay closer to the
%% actual signatures
tx_begin_response,          % Note that there is no dataflow from the AP
    tx_close_response,      % to the TM on many commands, so there is no
    tx_commit_response,     % need for a command type decl'n
    tx_information_command, tx_information_response,
    tx_open_response,
    tx_rollback_response: TYPE

ax_register_command, ax_register_response,
ax_unregister_command, ax_unregister_response,
xa_close_command, xa_close_response,
xa_commit_command, xa_commit_response,
xa_complete_command, xa_complete_response,
xa_end_command, xa_end_response,
xa_forget_command, xa_forget_response,
xa_open_command, xa_open_response,
xa_prepare_command, xa_prepare_response,
xa_recover_command, xa_recover_response,
xa_rollback_command, xa_rollback_response,
xa_start_command, xa_start_response: TYPE

q_type: {i: NAT | i < n} --> {t: TYPE | t < ar_requests}
r_type: {i: NAT | i < n} --> {t: TYPE | t < ar_resources}

COMPONENTS

ap: TYPE <= Function [<< ap_in1(i): r_type(i) |(i: NAT) i < n >>,
                      ap_begin_response_in: tx_begin_response,
                      ap_close_response_in: tx_close_response,
                      ap_commit_response_in: tx_commit_response,
                      ap_information_response_in: tx_information_response,
                      ap_open_response_in: tx_open_response,
                      ap_rollback_response_in: tx_rollback_response
                     -> << ap_out1(i): q_type(i) |(i: NAT) i < n >>,
                      ap_information_command_out: tx_information_command]

rm: TYPE <= { p: Function[rm_in1: qt,
                        rm_register_in: ax_register_response,
                        rm_unregister_in: ax_unregister_response,
                        rm_close_in: xa_close_command,
                        rm_commit_in: xa_commit_command,
                        rm_complete_in: xa_complete_command,
                        rm_end_in: xa_end_command,
                        rm_forget_in: xa_forget_command,
                        rm_open_in: xa_open_command,
                        rm_prepare_in: xa_prepare_command,
                        rm_recover_in: xa_recover_command,

```

```

        rm_rollback_in: xa_rollback_command,
        rm_start_in: xa_start_command
-> rm_out1: rt,
        rm_register_out: ax_register_command,
        rm_unregister_out: ax_unregister_command,
        rm_close_out: xa_close_response,
        rm_commit_out: xa_commit_response,
        rm_complete_out: xa_complete_response,
        rm_end_out: xa_end_response,
        rm_forget_out: xa_forget_response,
        rm_open_out: xa_open_response,
        rm_prepare_out: xa_prepare_response,
        rm_recover_out: xa_recover_response,
        rm_rollback_out: xa_rollback_response,
        rm_start_out: xa_start_response]
| qt < ar_requests AND rt < ar_resources }

rms: TYPE <= ARCHITECTURE [ << rm_in1(i): q_type(i) |(i: NAT) i < n >>,
    << rm_register_in(i): ax_register_response
        |(i: NAT) i < n >>,
    << rm_unregister_in(i): ax_unregister_response
        |(i: NAT) i < n >>,
    << rm_close_in(i): xa_close_command
        |(i: NAT) i < n >>,
    << rm_commit_in(i): xa_commit_command
        |(i: NAT) i < n >>,
    << rm_complete_in(i): xa_complete_command
        |(i: NAT) i < n >>,
    << rm_end_in(i): xa_end_command
        |(i: NAT) i < n >>,
    << rm_forget_in(i): xa_forget_command
        |(i: NAT) i < n >>,
    << rm_open_in(i): xa_open_command
        |(i: NAT) i < n >>,
    << rm_prepare_in(i): xa_prepare_command
        |(i: NAT) i < n >>,
    << rm_recover_in(i): xa_recover_command
        |(i: NAT) i < n >>,
    << rm_rollback_in(i): xa_rollback_command
        |(i: NAT) i < n >>,
    << rm_start_in(i): xa_start_command
        |(i: NAT) i < n >>
-> << rm_out1(i): r_type(i) |(i: NAT) i < n >>,
    << rm_register_out(i): ax_register_command
        |(i: NAT) i < n >>,
    << rm_unregister_out(i): ax_unregister_command
        |(i: NAT) i < n >>,
    << rm_close_out(i): xa_close_response
        |(i: NAT) i < n >>,
    << rm_commit_out(i): xa_commit_response
        |(i: NAT) i < n >>,
    << rm_complete_out(i): xa_complete_response
        |(i: NAT) i < n >>,
    << rm_end_out(i): xa_end_response
        |(i: NAT) i < n >>,
    << rm_forget_out(i): xa_forget_response
        |(i: NAT) i < n >>,
    << rm_open_out(i): xa_open_response
        |(i: NAT) i < n >>,
    << rm_prepare_out(i): xa_prepare_response
        |(i: NAT) i < n >>,
    << rm_recover_out(i): xa_recover_response
        |(i: NAT) i < n >>,
    << rm_rollback_out(i): xa_rollback_response

```

```

        |(i: NAT) i < n >,
<< rm_start_out(i): xa_start_response
        |(i: NAT) i < n >>]

tm: TYPE <= Function [tm_information_command_in: tx_information_command,
    << tm_register_in(i): ax_register_command
        |(i: NAT) i < n >,
    << tm_unregister_in(i): ax_unregister_command
        |(i: NAT) i < n >,
    << tm_close_in(i): xa_close_response
        |(i: NAT) i < n >,
    << tm_commit_in(i): xa_commit_response
        |(i: NAT) i < n >,
    << tm_complete_in(i): xa_complete_response
        |(i: NAT) i < n >,
    << tm_end_in(i): xa_end_response
        |(i: NAT) i < n >,
    << tm_forget_in(i): xa_forget_response
        |(i: NAT) i < n >,
    << tm_open_in(i): xa_open_response
        |(i: NAT) i < n >,
    << tm_prepare_in(i): xa_prepare_response
        |(i: NAT) i < n >,
    << tm_recover_in(i): xa_recover_response
        |(i: NAT) i < n >,
    << tm_rollback_in(i): xa_rollback_response
        |(i: NAT) i < n >,
    << tm_start_in(i): xa_start_response
        |(i: NAT) i < n >>
-> tm_begin_response_out: tx_begin_response,
    tm_close_response_out: tx_close_response,
    tm_commit_response_out: tx_commit_response,
    tm_information_response_out: tx_information_response,
    tm_open_response_out: tx_open_response,
    tm_rollback_response_out: tx_rollback_response,
    << tm_register_out(i): ax_register_response
        |(i: NAT) i < n >,
    << tm_unregister_out(i): ax_unregister_response
        |(i: NAT) i < n >,
    << tm_close_out(i): xa_close_command
        |(i: NAT) i < n >,
    << tm_commit_out(i): xa_commit_command
        |(i: NAT) i < n >,
    << tm_complete_out(i): xa_complete_command
        |(i: NAT) i < n >,
    << tm_end_out(i): xa_end_command
        |(i: NAT) i < n >,
    << tm_forget_out(i): xa_forget_command
        |(i: NAT) i < n >,
    << tm_open_out(i): xa_open_command
        |(i: NAT) i < n >,
    << tm_prepare_out(i): xa_prepare_command
        |(i: NAT) i < n >,
    << tm_recover_out(i): xa_recover_command
        |(i: NAT) i < n >,
    << tm_rollback_out(i): xa_rollback_command
        |(i: NAT) i < n >,
    << tm_start_out(i): xa_start_command
        |(i: NAT) i < n >>]

the_ap: ap
the_rms: rms
the_tm: tm

```

CONFIGURATION

```

rms_contents: CONSTRAINT =
    (FORALL y: COMPONENT)
        [y PROPERLY_CONTAINED_IN the_rms => (EXISTS z: rm) y CONTAINED_IN z]
rm_location: CONSTRAINT =
    (FORALL y: rm) y CONTAINED_IN the_rms

ar_1: CONNECTION =
    (FORALL i: NAT | i < n)
        (EXISTS c: Channel<q_type(i)>)
            Connects(c, the_ap.ap_out1(i), the_rms.rm_in1(i))
ar_2: CONNECTION =
    (FORALL i: NAT | i < n)
        (EXISTS c: Channel<r_type(i)>)
            Connects(c, the_rms.rm_out1(i), the_ap.ap_in1(i))

tx_1: CONNECTION =
    (EXISTS c: Channel<tx_information_command>)
        Connects(c, the_ap.ap_information_command_out,
                  the_tm.tm_information_command_in)
tx_2a: CONNECTION =
    (EXISTS c: Channel<tx_begin_response>)
        Connects(c, the_tm.tm_begin_response_out,
                  the_ap.ap_begin_response_in)
tx_2b: CONNECTION =
    (EXISTS c: Channel<tx_close_response>)
        Connects(c, the_tm.tm_close_response_out,
                  the_ap.ap_close_response_in)
tx_2c: CONNECTION =
    (EXISTS c: Channel<tx_commit_response>)
        Connects(c, the_tm.tm_commit_response_out,
                  the_ap.ap_commit_response_in)
tx_2d: CONNECTION =
    (EXISTS c: Channel<tx_information_response>)
        Connects(c, the_tm.tm_information_response_out,
                  the_ap.ap_information_response_in)
tx_2e: CONNECTION =
    (EXISTS c: Channel<tx_open_response>)
        Connects(c, the_tm.tm_open_response_out,
                  the_ap.ap_open_response_in)
tx_2f: CONNECTION =
    (EXISTS c: Channel<tx_rollback_response>)
        Connects(c, the_tm.tm_rollback_response_out,
                  the_ap.ap_rollback_response_in)

xa_1a: CONNECTION =
    (FORALL i: NAT | i < n)
        (EXISTS c: Channel<ax_register_response>)
            Connects(c, the_tm.tm_register_response_out(i),
                      the_rms.rm_register_response_in(i))
xa_1b: CONNECTION =
    (FORALL i: NAT | i < n)
        (EXISTS c: Channel<ax_unregister_response>)
            Connects(c, the_tm.tm_unregister_response_out(i),
                      the_rms.rm_unregister_response_in(i))
xa_1c: CONNECTION =
    (FORALL i: NAT | i < n)
        (EXISTS c: Channel<xa_close_command>)
            Connects(c, the_tm.tm_close_command_out(i),
                      the_rms.rm_close_command_in(i))
xa_1d: CONNECTION =
    (FORALL i: NAT | i < n)
        (EXISTS c: Channel<xa_commit_command>)

```

```

        Connects(c, the_tm.tm_commit_command_out(i),
                  the_rms.rm_commit_command_in(i))

xa_1e: CONNECTION =
    (FORALL i: NAT | i < n)
        (EXISTS c: Channel<xa_complete_command>)
            Connects(c, the_tm.tm_complete_command_out(i),
                      the_rms.rm_complete_command_in(i))

xa_1f: CONNECTION =
    (FORALL i: NAT | i < n)
        (EXISTS c: Channel<xa_end_command>)
            Connects(c, the_tm.tm_end_command_out(i),
                      the_rms.rm_end_command_in(i))

xa_1g: CONNECTION =
    (FORALL i: NAT | i < n)
        (EXISTS c: Channel<xa_forget_command>)
            Connects(c, the_tm.tm_forget_command_out(i),
                      the_rms.rm_forget_command_in(i))

xa_1h: CONNECTION =
    (FORALL i: NAT | i < n)
        (EXISTS c: Channel<xa_open_command>)
            Connects(c, the_tm.tm_open_command_out(i),
                      the_rms.rm_open_command_in(i))

xa_1i: CONNECTION =
    (FORALL i: NAT | i < n)
        (EXISTS c: Channel<xa_prepare_command>)
            Connects(c, the_tm.tm_prepare_command_out(i),
                      the_rms.rm_prepare_command_in(i))

xa_1j: CONNECTION =
    (FORALL i: NAT | i < n)
        (EXISTS c: Channel<xa_recover_command>)
            Connects(c, the_tm.tm_recover_command_out(i),
                      the_rms.rm_recover_command_in(i))

xa_1k: CONNECTION =
    (FORALL i: NAT | i < n)
        (EXISTS c: Channel<xa_rollback_command>)
            Connects(c, the_tm.tm_rollback_command_out(i),
                      the_rms.rm_rollback_command_in(i))

xa_1l: CONNECTION =
    (FORALL i: NAT | i < n)
        (EXISTS c: Channel<xa_start_command>)
            Connects(c, the_tm.tm_start_command_out(i),
                      the_rms.rm_start_command_in(i))

xa_2a: CONNECTION =
    (FORALL i: NAT | i < n)
        (EXISTS c: Channel<ax_register_command>)
            Connects(c, the_rms.rm_register_command_out(i),
                      the_tm.tm_register_command_in(i))

xa_2b: CONNECTION =
    (FORALL i: NAT | i < n)
        (EXISTS c: Channel<ax_unregister_command>)
            Connects(c, the_rms.rm_unregister_command_out(i),
                      the_tm.tm_unregister_command_in(i))

xa_2c: CONNECTION =
    (FORALL i: NAT | i < n)
        (EXISTS c: Channel<xa_close_response>)
            Connects(c, the_rms.rm_close_response_out(i),
                      the_tm.tm_close_response_in(i))

xa_2d: CONNECTION =
    (FORALL i: NAT | i < n)
        (EXISTS c: Channel<xa_commit_response>)
            Connects(c, the_rms.rm_commit_response_out(i),
                      the_tm.tm_commit_response_in(i))

xa_2e: CONNECTION =
    (FORALL i: NAT | i < n)

```

```

        (EXISTS c: Channel<xa_complete_response>)
            Connects(c, the_rms.rm_complete_response_out(i),
                      the_tm.tm_complete_response_in(i))

xa_2f: CONNECTION =
        (FORALL i: NAT | i < n)
            (EXISTS c: Channel<xa_end_response>)
                Connects(c, the_rms.rm_end_response_out(i),
                          the_tm.tm_end_response_in(i))

xa_2g: CONNECTION =
        (FORALL i: NAT | i < n)
            (EXISTS c: Channel<xa_forget_response>)
                Connects(c, the_rms.rm_forget_response_out(i),
                          the_tm.tm_forget_response_in(i))

xa_2h: CONNECTION =
        (FORALL i: NAT | i < n)
            (EXISTS c: Channel<xa_open_response>)
                Connects(c, the_rms.rm_open_response_out(i),
                          the_tm.tm_open_response_in(i))

xa_2i: CONNECTION =
        (FORALL i: NAT | i < n)
            (EXISTS c: Channel<xa_prepare_response>)
                Connects(c, the_rms.rm_prepare_response_out(i),
                          the_tm.tm_prepare_response_in(i))

xa_2j: CONNECTION =
        (FORALL i: NAT | i < n)
            (EXISTS c: Channel<xa_recover_response>)
                Connects(c, the_rms.rm_recover_response_out(i),
                          the_tm.tm_recover_response_in(i))

xa_2k: CONNECTION =
        (FORALL i: NAT | i < n)
            (EXISTS c: Channel<xa_rollback_response>)
                Connects(c, the_rms.rm_rollback_response_out(i),
                          the_tm.tm_rollback_response_in(i))

xa_2l: CONNECTION =
        (FORALL i: NAT | i < n)
            (EXISTS c: Channel<xa_start_response>)
                Connects(c, the_rms.rm_start_response_out(i),
                          the_tm.tm_start_response_in(i))

END x_open_concrete_df

```

```
%%% Introduce the actual types on the TX and XA interfaces.  
%%% Note that this is the first use of X/Open_style: we need the type  
%%% definitions.
```

```
x_open_truetypes_df: ARCHITECTURE [ -> ]  
  
IMPORTING ALL FROM Dataflow_style, X_Open_style  
  
BEGIN  
  
n: NAT % Number of resource managers, a parameter in the specification  
  
ar_requests, ar_resources: TYPE  
  
q_type: {i: NAT | i < n} --> {t: TYPE | t < ar_requests}  
r_type: {i: NAT | i < n} --> {t: TYPE | t < ar_resources}  
  
COMPONENTS  
  
ap: TYPE <= Function [<< ap_in1(i): r_type(i) |(i: NAT) i < n >>,  
                      ap_begin_response_in: INT,  
                      ap_close_response_in: INT,  
                      ap_commit_response_in: INT,  
                      ap_information_response_in: INT,  
                      ap_open_response_in: INT,  
                      ap_rollback_response_in: INT  
                    -> << ap_out1(i): q_type(i) |(i: NAT) i < n >>,  
                      ap_information_command_out: TX_Info]  
  
rm: TYPE <= { p: Function[rm_in1: qt,  
                           rm_register_in: INT,  
                           rm_unregister_in: INT,  
                           rm_close_in: XA_Info X INT^2,  
                           rm_commit_in: X_Id X INT^2,  
                           rm_complete_in: INT^4,  
                           rm_end_in: X_id X INT^2,  
                           rm_forget_in: X_Id X INT^2,  
                           rm_open_in: XA_Info X INT^2,  
                           rm_prepare_in: X_Id X INT^2,  
                           rm_recover_in: X_Ids X INT^3,  
                           rm_rollback_in: X_Id X INT^2,  
                           rm_start_in: X_Id X INT^2  
                         -> rm_out1: rt,  
                           rm_register_out: X_Id X INT^2,  
                           rm_unregister_out: INT^2,  
                           rm_close_out: INT,  
                           rm_commit_out: INT,  
                           rm_complete_out: INT,  
                           rm_end_out: INT,  
                           rm_forget_out: INT,  
                           rm_open_out: INT,  
                           rm_prepare_out: INT,  
                           rm_recover_out: INT,  
                           rm_rollback_out: INT,  
                           rm_start_out: INT]  
                         | qt < ar_requests AND rt < ar_resources }  
  
rms: TYPE <= ARCHITECTURE [<< rm_in1(i): q_type(i) |(i: NAT) i < n >>,  
                            << rm_register_in(i): INT  
                               |(i: NAT) i < n >>,  
                            << rm_unregister_in(i): INT  
                               |(i: NAT) i < n >>,  
                            << rm_close_in(i): XA_Info X INT^2
```

```

        |(i: NAT) i < n >>,
<< rm_commit_in(i): X_Id X INT^2
        |(i: NAT) i < n >>,
<< rm_complete_in(i): INT^4
        |(i: NAT) i < n >>,
<< rm_end_in(i): X_id X INT^2
        |(i: NAT) i < n >>,
<< rm_forget_in(i): X_Id X INT^2
        |(i: NAT) i < n >>,
<< rm_open_in(i): XA_Info X INT^2
        |(i: NAT) i < n >>,
<< rm_prepare_in(i): X_Id X INT^2
        |(i: NAT) i < n >>,
<< rm_recover_in(i): X_Ids X INT^3
        |(i: NAT) i < n >>,
<< rm_rollback_in(i): X_Id X INT^2
        |(i: NAT) i < n >>,
<< rm_start_in(i): X_Id X INT^2
        |(i: NAT) i < n >>
-> << rm_out1(i): r_type(i) |(i: NAT) i < n >>,
    << rm_register_out(i): X_Id X INT^2
        |(i: NAT) i < n >>,
    << rm_unregister_out(i): INT^2
        |(i: NAT) i < n >>,
    << rm_close_out(i): INT
        |(i: NAT) i < n >>,
    << rm_commit_out(i): INT
        |(i: NAT) i < n >>,
    << rm_complete_out(i): INT
        |(i: NAT) i < n >>,
    << rm_end_out(i): INT
        |(i: NAT) i < n >>,
    << rm_forget_out(i): INT
        |(i: NAT) i < n >>,
    << rm_open_out(i): INT
        |(i: NAT) i < n >>,
    << rm_prepare_out(i): INT
        |(i: NAT) i < n >>,
    << rm_recover_out(i): INT
        |(i: NAT) i < n >>,
    << rm_rollback_out(i): INT
        |(i: NAT) i < n >>,
    << rm_start_out(i): INT
        |(i: NAT) i < n >>]

```

```

tm: TYPE <= Function [tm_information_command_in: TX_Info,
    << tm_register_in(i): X_Id X INT^2
        |(i: NAT) i < n >>,
    << tm_unregister_in(i): INT^2
        |(i: NAT) i < n >>,
    << tm_close_in(i): INT
        |(i: NAT) i < n >>,
    << tm_commit_in(i): INT
        |(i: NAT) i < n >>,
    << tm_complete_in(i): INT
        |(i: NAT) i < n >>,
    << tm_end_in(i): INT
        |(i: NAT) i < n >>,
    << tm_forget_in(i): INT
        |(i: NAT) i < n >>,
    << tm_open_in(i): INT
        |(i: NAT) i < n >>,
    << tm_prepare_in(i): INT
        |(i: NAT) i < n >>,

```

```

        << tm_recover_in(i): INT
            |(i: NAT) i < n >>,
        << tm_rollback_in(i): INT
            |(i: NAT) i < n >>,
        << tm_start_in(i): INT
            |(i: NAT) i < n >>
-> tm_begin_response_out: INT,
    tm_close_response_out: INT,
    tm_commit_response_out: INT,
    tm_information_response_out: INT,
    tm_open_response_out: INT,
    tm_rollback_response_out: INT,
    << tm_register_out(i): INT
        |(i: NAT) i < n >>,
    << tm_unregister_out(i): INT
        |(i: NAT) i < n >>,
    << tm_close_out(i): XA_Info X INT^2
        |(i: NAT) i < n >>,
    << tm_commit_out(i): X_Id X INT^2
        |(i: NAT) i < n >>,
    << tm_complete_out(i): INT^4
        |(i: NAT) i < n >>,
    << tm_end_out(i): X_id X INT^2
        |(i: NAT) i < n >>,
    << tm_forget_out(i): X_Id X INT^2
        |(i: NAT) i < n >>,
    << tm_open_out(i): XA_Info X INT^2
        |(i: NAT) i < n >>,
    << tm_prepare_out(i): X_Id X INT^2
        |(i: NAT) i < n >>,
    << tm_recover_out(i): X_Ids X INT^3
        |(i: NAT) i < n >>,
    << tm_rollback_out(i): X_Id X INT^2
        |(i: NAT) i < n >>,
    << tm_start_out(i): X_Id X INT^2
        |(i: NAT) i < n >>]

```

the_ap: ap
the_rms: rms
the_tm: tm

CONFIGURATION

```

rms_contents: CONSTRAINT =
    (FORALL y: COMPONENT)
        [y PROPERLY_CONTAINED_IN the_rms => (EXISTS z: rm) y CONTAINED_IN z]
rm_location: CONSTRAINT =
    (FORALL y: rm) y CONTAINED_IN the_rms

ar_1: CONNECTION =
    (FORALL i: NAT | i < n)
        (EXISTS c: Channel<q_type(i)>)
            Connects(c, the_ap.ap_out1(i), the_rms.rm_in1(i))
ar_2: CONNECTION =
    (FORALL i: NAT | i < n)
        (EXISTS c: Channel<r_type(i)>)
            Connects(c, the_rms.rm_out1(i), the_ap.ap_in1(i))

tx_1: CONNECTION =
    (EXISTS c: Channel<TX_Info>)
        Connects(c, the_ap.ap_information_command_out,
                  the_tm.tm_information_command_in)

tx_2a: CONNECTION =
    (EXISTS c: Channel<INT>)

```

```

        Connects(c, the_tm.tm_begin_response_out,
                  the_ap.ap_begin_response_in)
tx_2b: CONNECTION =
        (EXISTS c: Channel<INT>)
        Connects(c, the_tm.tm_close_response_out,
                  the_ap.ap_close_response_in)
tx_2c: CONNECTION =
        (EXISTS c: Channel<INT>)
        Connects(c, the_tm.tm_commit_response_out,
                  the_ap.ap_commit_response_in)
tx_2d: CONNECTION =
        (EXISTS c: Channel<INT>)
        Connects(c, the_tm.tm_information_response_out,
                  the_ap.ap_information_response_in)
tx_2e: CONNECTION =
        (EXISTS c: Channel<INT>)
        Connects(c, the_tm.tm_open_response_out,
                  the_ap.ap_open_response_in)
tx_2f: CONNECTION =
        (EXISTS c: Channel<INT>)
        Connects(c, the_tm.tm_rollback_response_out,
                  the_ap.ap_rollback_response_in)

xa_1a: CONNECTION =
        (FORALL i: NAT | i < n)
        (EXISTS c: Channel<INT>)
        Connects(c, the_tm.tm_register_response_out(i),
                  the_rms.rm_register_response_in(i))
xa_1b: CONNECTION =
        (FORALL i: NAT | i < n)
        (EXISTS c: Channel<INT>)
        Connects(c, the_tm.tm_unregister_response_out(i),
                  the_rms.rm_unregister_response_in(i))
xa_1c: CONNECTION =
        (FORALL i: NAT | i < n)
        (EXISTS c: Channel<XA_Info X INT^2>)
        Connects(c, the_tm.tm_close_command_out(i),
                  the_rms.rm_close_command_in(i))
xa_1d: CONNECTION =
        (FORALL i: NAT | i < n)
        (EXISTS c: Channel<X_Id X INT^2>)
        Connects(c, the_tm.tm_commit_command_out(i),
                  the_rms.rm_commit_command_in(i))
xa_1e: CONNECTION =
        (FORALL i: NAT | i < n)
        (EXISTS c: Channel<INT^4>)
        Connects(c, the_tm.tm_complete_command_out(i),
                  the_rms.rm_complete_command_in(i))
xa_1f: CONNECTION =
        (FORALL i: NAT | i < n)
        (EXISTS c: Channel<X_id X INT^2>)
        Connects(c, the_tm.tm_end_command_out(i),
                  the_rms.rm_end_command_in(i))
xa_1g: CONNECTION =
        (FORALL i: NAT | i < n)
        (EXISTS c: Channel<X_Id X INT^2>)
        Connects(c, the_tm.tm_forget_command_out(i),
                  the_rms.rm_forget_command_in(i))
xa_1h: CONNECTION =
        (FORALL i: NAT | i < n)
        (EXISTS c: Channel<XA_Info X INT^2>)
        Connects(c, the_tm.tm_open_command_out(i),
                  the_rms.rm_open_command_in(i))
xa_1i: CONNECTION =

```

```

(FORALL i: NAT | i < n)
  (EXISTS c: Channel<X_Id X INT^2>)
    Connects(c, the_tm.tm_prepare_command_out(i),
              the_rms.rm_prepare_command_in(i))

xa_1j: CONNECTION =
  (FORALL i: NAT | i < n)
    (EXISTS c: Channel<X_Ids X INT^3>)
      Connects(c, the_tm.tm_recover_command_out(i),
                the_rms.rm_recover_command_in(i))

xa_1k: CONNECTION =
  (FORALL i: NAT | i < n)
    (EXISTS c: Channel<X_Id X INT^2>)
      Connects(c, the_tm.tm_rollback_command_out(i),
                the_rms.rm_rollback_command_in(i))

xa_1l: CONNECTION =
  (FORALL i: NAT | i < n)
    (EXISTS c: Channel<X_Id X INT^2>)
      Connects(c, the_tm.tm_start_command_out(i),
                the_rms.rm_start_command_in(i))

xa_2a: CONNECTION =
  (FORALL i: NAT | i < n)
    (EXISTS c: Channel<X_Id X INT^2>)
      Connects(c, the_rms.rm_register_command_out(i),
                the_tm.tm_register_command_in(i))

xa_2b: CONNECTION =
  (FORALL i: NAT | i < n)
    (EXISTS c: Channel<INT^2>)
      Connects(c, the_rms.rm_unregister_command_out(i),
                the_tm.tm_unregister_command_in(i))

xa_2c: CONNECTION =
  (FORALL i: NAT | i < n)
    (EXISTS c: Channel<INT>)
      Connects(c, the_rms.rm_close_response_out(i),
                the_tm.tm_close_response_in(i))

xa_2d: CONNECTION =
  (FORALL i: NAT | i < n)
    (EXISTS c: Channel<INT>)
      Connects(c, the_rms.rm_commit_response_out(i),
                the_tm.tm_commit_response_in(i))

xa_2e: CONNECTION =
  (FORALL i: NAT | i < n)
    (EXISTS c: Channel<INT>)
      Connects(c, the_rms.rm_complete_response_out(i),
                the_tm.tm_complete_response_in(i))

xa_2f: CONNECTION =
  (FORALL i: NAT | i < n)
    (EXISTS c: Channel<INT>)
      Connects(c, the_rms.rm_end_response_out(i),
                the_tm.tm_end_response_in(i))

xa_2g: CONNECTION =
  (FORALL i: NAT | i < n)
    (EXISTS c: Channel<INT>)
      Connects(c, the_rms.rm_forget_response_out(i),
                the_tm.tm_forget_response_in(i))

xa_2h: CONNECTION =
  (FORALL i: NAT | i < n)
    (EXISTS c: Channel<INT>)
      Connects(c, the_rms.rm_open_response_out(i),
                the_tm.tm_open_response_in(i))

xa_2i: CONNECTION =
  (FORALL i: NAT | i < n)
    (EXISTS c: Channel<INT>)
      Connects(c, the_rms.rm_prepare_response_out(i),
                the_tm.tm_prepare_response_in(i))

```

```
xa_2j: CONNECTION =
  (FORALL i: NAT | i < n)
    (EXISTS c: Channel<INT>)
      Connects(c, the_rms.rm_recover_response_out(i),
                the_tm.tm_recover_response_in(i))

xa_2k: CONNECTION =
  (FORALL i: NAT | i < n)
    (EXISTS c: Channel<INT>)
      Connects(c, the_rms.rm_rollback_response_out(i),
                the_tm.tm_rollback_response_in(i))

xa_2l: CONNECTION =
  (FORALL i: NAT | i < n)
    (EXISTS c: Channel<INT>)
      Connects(c, the_rms.rm_start_response_out(i),
                the_tm.tm_start_response_in(i))

END x_open_truetype_df
```

```

%%% Replace TX and XA dataflow by procedure calls, using
%%% pre-defined procedure call varieties. This requires first replacing
%%% Functions by ARCHITECTURES, so that the procedure declarations can be stuck
%%% in the right places. (AP is changed to a ARCHITECTURE for uniformity, and
%%% to eliminate the dependence on dataflow style.)

x_open_semiproc: ARCHITECTURE [ -> ]
  IMPORTING ALL FROM Dataflow_style,
    X_Open_style % defines XA_Close_Procedure, ..., TX_Begin_Procedure,
...
BEGIN
  n: NAT % Number of resource managers, a parameter in the specification
  ar_requests, ar_resources: TYPE
  q_type: {i: NAT | i < n} --> {t: TYPE | t < ar_requests}
  r_type: {i: NAT | i < n} --> {t: TYPE | t < ar_resources}

COMPONENTS
  ap: TYPE <= ARCHITECTURE [<< ap_in1(i): r_type(i) |(i: NAT) i < n >>
    -> << ap_out1(i): q_type(i) |(i: NAT) i < n >>]

  rm: TYPE <= { m: ARCHITECTURE [rm_in1: qt -> rm_out1: rt]
    EXPORTING ALL
    BEGIN
      close: XA_Close_Procedure
        [info: XA_Info, rmid: INT, flags: INT
         -> ret: INT]
      commit: XA_Commit_Procedure
        [id: X_Id, rmid: INT, flags: INT
         -> ret: INT]
      complete: XA_Complete_Procedure
        [hndl: INT, retval: INT,
         rmid: INT, flags: INT
         -> ret: INT]
      end: XA_End_Procedure
        [id: X_Id, rmid: INT, flags: INT
         -> ret: INT]
      forget: XA_Forget_Procedure
        [id: X_Id, rmid: INT, flags: INT
         -> ret: INT]
      open: XA_Open_Procedure
        [info: XA_Info, rmid: INT, flags: INT
         -> ret: INT]
      prepare: XA_Prepares_Procedure
        [id: X_Id, rmid: INT, flags: INT
         -> ret: INT]
      recover: XA_Recover_Procedure
        [ids: X_Ids, count: INT,
         rmid: INT, flags: INT
         -> ret: INT]
      rollback: XA_Rollback_Procedure
        [id: X_Id, rmid: INT, flags: INT
         -> ret: INT]
      start: XA_Start_Procedure
        [id: X_Id, rmid: INT, flags: INT
         -> ret: INT]
    END m
    | qt < ar_requests AND rt < ar_resources }

```

```

rms: TYPE <= ARCHITECTURE [<< rm_in1(i): q_type(i) |(i: NAT) i < n >>
                           -> << rm_out1(i): r_type(i) |(i: NAT) i < n >>]

tm: TYPE <= ARCHITECTURE [ -> ]
    EXPORTING ALL
    BEGIN
        register: AX_Register_Procedure
            [id: X_Id, rmid: INT, flags: INT
             -> ret: INT]
        unregister: AX_Unregister_Procedure
            [rmid: INT, flags: INT
             -> ret: INT]
        begin: TX_Begin_Procedure [ -> ret: INT]
        close: TX_Close_Procedure [ -> ret: INT]
        commit: TX_Commit_Procedure [ -> ret: INT]
        information: TX_Info_Procedure [info: TX_Info -> ret: INT]
        open: TX_Open_Procedure [ -> ret: INT]
        rollback: TX_Rollback_Procedure [ -> ret: INT]
    END tm

the_ap: ap
the_rms: rms
the_tm: tm

CONFIGURATION

rms_contents: CONSTRAINT =
    (FORALL y: COMPONENT)
        [y PROPERLY_CONTAINED_IN the_rms => (EXISTS z: rm) y CONTAINED_IN z]
rm_location: CONSTRAINT =
    (FORALL y: rm) y CONTAINED_IN the_rms

ar_1: CONNECTION =
    (FORALL i: NAT | i < n)
        (EXISTS c: Channel<q_type(i)>)
            Connects(c, the_ap.ap_out1(i), the_rms.rm_in1(i))
ar_2: CONNECTION =
    (FORALL i: NAT | i < n)
        (EXISTS c: Channel<r_type(i)>)
            Connects(c, the_rms.rm_out1(i), the_ap.ap_in1(i))

%% For now, let's make these a bit more readable by (implicitly)
%% existentially quantifying the call sites away. (Of course, we'll
%% eventually need them in the mapping, but mappings can be hidden
%% behind the scenes on the transformational approach.)

tx: CONSTRAINT =
    Called_From(the_tm.begin, the_ap)
    AND Called_From(the_tm.close, the_ap)
    AND Called_From(the_tm.commit, the_ap)
    AND Called_From(the_tm.information, the_ap)
    AND Called_From(the_tm.open, the_ap)
    AND Called_From(the_tm.rollback, the_ap)

xa: CONSTRAINT =
    (FORALL y: rm)
        [Called_From(the_tm.register, y)
         AND Called_From(the_tm.unregister, y)
         AND Called_From(y.close, the_tm)
         AND Called_From(y.commit, the_tm)
         AND Called_From(y.complete, the_tm)
         AND Called_From(y.end, the_tm)
         AND Called_From(y.forget, the_tm)
         AND Called_From(y.open, the_tm)]

```

```
AND Called_From(y.prepare, the_tm)
AND Called_From(y.recover, the_tm)
AND Called_From(y.rollback, the_tm)
AND Called_From(y.start, the_tm)]  
END x_open_semiproc
```

```

%%% Replace AR dataflow by a remote procedure call, Note that use of
%%% Dataflow_style has been completely eliminated.

x_open_proc_1: ARCHITECTURE [ -> ]

    IMPORTING ALL FROM X_Open_style,
        RPC_style      % Introduces Remotely Callable Procedures,
                        %     RPCs, an implementation of PROCEDURES

BEGIN

COMPONENTS

    ap: TYPE <= ARCHITECTURE [ -> ]

    rm: TYPE <= { m: ARCHITECTURE [ -> ]
                  EXPORTING ALL
                  BEGIN
                      access_function: RPC [in: qt -> out: rt]
                      close: XA_Close_Procedure
                          [info: XA_Info, rmid: INT, flags: INT
                           -> ret: INT]
                      commit: XA_Commit_Procedure
                          [id: X_Id, rmid: INT, flags: INT
                           -> ret: INT]
                      complete: XA_Complete_Procedure
                          [hdl: INT, retval: INT,
                           rmid: INT, flags: INT
                           -> ret: INT]
                      end: XA_End_Procedure
                          [id: X_Id, rmid: INT, flags: INT
                           -> ret: INT]
                      forget: XA_Forget_Procedure
                          [id: X_Id, rmid: INT, flags: INT
                           -> ret: INT]
                      open: XA_Open_Procedure
                          [info: XA_Info, rmid: INT, flags: INT
                           -> ret: INT]
                      prepare: XA_Prepares_Procedure
                          [id: X_Id, rmid: INT, flags: INT
                           -> ret: INT]
                      recover: XA_Recover_Procedure
                          [ids: X_Ids, count: INT,
                           rmid: INT, flags: INT
                           -> ret: INT]
                      rollback: XA_Rollback_Procedure
                          [id: X_Id, rmid: INT, flags: INT
                           -> ret: INT]
                      start: XA_Start_Procedure
                          [id: X_Id, rmid: INT, flags: INT
                           -> ret: INT]
                  END m
                  | qt < ar_requests AND rt < ar_resources }

    rms: TYPE <= ARCHITECTURE [ -> ]

    tm: TYPE <= ARCHITECTURE [ -> ]
                  EXPORTING ALL
                  BEGIN
                      register: AX_Register_Procedure
                          [id: X_Id, rmid: INT, flags: INT
                           -> ret: INT]
                      unregister: AX_Unregister_Procedure
                          [rmid: INT, flags: INT

```

```

        -> ret: INT]
begin: TX_Begin_Procedure [ -> ret: INT]
close: TX_Close_Procedure [ -> ret: INT]
commit: TX_Commit_Procedure [ -> ret: INT]
information: TX_Info_Procedure [info: TX_Info -> ret: INT]
open: TX_Open_Procedure [ -> ret: INT]
rollback: TX_Rollback_Procedure [ -> ret: INT]
END tm

the_ap: ap
the_rms: rms
the_tm: tm

CONFIGURATION

rms_contents: CONSTRAINT =
  (FORALL y: COMPONENT)
    [y PROPERLY_CONTAINED_IN the_rms => (EXISTS z: rm) y CONTAINED_IN z]
rm_location: CONSTRAINT =
  (FORALL y: rm) y CONTAINED_IN the_rms

ar: CONSTRAINT =
  (FORALL y: rm) Called_From(y.access_function, the_ap)

tx: CONSTRAINT =
  Called_From(the_tm.begin, the_ap)
  AND Called_From(the_tm.close, the_ap)
  AND Called_From(the_tm.commit, the_ap)
  AND Called_From(the_tm.information, the_ap)
  AND Called_From(the_tm.open, the_ap)
  AND Called_From(the_tm.rollback, the_ap)

xa: CONSTRAINT =
  (FORALL y: rm)
    [Called_From(the_tm.register, y)
     AND Called_From(the_tm.unregister, y)
     AND Called_From(y.close, the_tm)
     AND Called_From(y.commit, the_tm)
     AND Called_From(y.complete, the_tm)
     AND Called_From(y.end, the_tm)
     AND Called_From(y.forget, the_tm)
     AND Called_From(y.open, the_tm)
     AND Called_From(y.prepare, the_tm)
     AND Called_From(y.recover, the_tm)
     AND Called_From(y.rollback, the_tm)
     AND Called_From(y.start, the_tm)]


END x_open_proc_1

```

```

%%% Replace AR dataflow by a pair of remote procedure calls. Synchronize,
%%% but don't block waiting for slow resource managers.

x_open_proc_2: ARCHITECTURE [ -> ]

IMPORTING ALL FROM X_Open_style,
RPC_style           % Introduces Remotely Callable Procedures,
                     % RPCs, an implementation of PROCEDUREs

BEGIN

n: NAT    % Number of resource managers, a parameter in the specification

ar_requests, ar_resources: TYPE

COMPONENTS

ap: TYPE <= ARCHITECTURE [ -> ]
      EXPORTING ALL
      BEGIN
          { return_resource(i): RPC [in: r_type(i) -> ] |(i: NAT) i < n }
      END ap

rm: TYPE <= { m: ARCHITECTURE [ -> ]
      EXPORTING ALL
      BEGIN
          request_resource: RPC [in: qt -> ]
          close: XA_Close_Procedure
              [info: XA_Info, rmid: INT, flags: INT
               -> ret: INT]
          commit: XA_Commit_Procedure
              [id: X_Id, rmid: INT, flags: INT
               -> ret: INT]
          complete: XA_Complete_Procedure
              [hdl: INT, retval: INT,
               rmid: INT, flags: INT
               -> ret: INT]
          end: XA_End_Procedure
              [id: X_Id, rmid: INT, flags: INT
               -> ret: INT]
          forget: XA_Forget_Procedure
              [id: X_Id, rmid: INT, flags: INT
               -> ret: INT]
          open: XA_Open_Procedure
              [info: XA_Info, rmid: INT, flags: INT
               -> ret: INT]
          prepare: XA_Prepares_Procedure
              [id: X_Id, rmid: INT, flags: INT
               -> ret: INT]
          recover: XA_Recover_Procedure
              [ids: X_Ids, count: INT,
               rmid: INT, flags: INT
               -> ret: INT]
          rollback: XA_Rollback_Procedure
              [id: X_Id, rmid: INT, flags: INT
               -> ret: INT]
          start: XA_Start_Procedure
              [id: X_Id, rmid: INT, flags: INT
               -> ret: INT]
      END m
      | qt < ar_requests AND rt < ar_resources }

rms: TYPE <= ARCHITECTURE [ -> ]

```

```

tm: TYPE <= ARCHITECTURE [ -> ]
    EXPORTING ALL
    BEGIN
        register: AX_Register_Procedure
            [id: X_Id, rmid: INT, flags: INT
             -> ret: INT]
        unregister: AX_Unregister_Procedure
            [rmid: INT, flags: INT
             -> ret: INT]
        begin: TX_Begin_Procedure [ -> ret: INT]
        close: TX_Close_Procedure [ -> ret: INT]
        commit: TX_Commit_Procedure [ -> ret: INT]
        information: TX_Info_Procedure [info: TX_Info -> ret: INT]
        open: TX_Open_Procedure [ -> ret: INT]
        rollback: TX_Rollback_Procedure [ -> ret: INT]
    END tm

```

```

the_ap: ap
the_rms: rms
the_tm: tm

```

CONFIGURATION

```

rms_contents: CONSTRAINT =
    (FORALL y: COMPONENT)
        [y PROPERLY_CONTAINED_IN the_rms => (EXISTS z: rm) y CONTAINED_IN z]
rm_location: CONSTRAINT =
    (FORALL y: rm) y CONTAINED_IN the_rms

ar_1: CONSTRAINT =
    (FORALL y: rm) Called_From(y.request_resource, the_ap)

ar_2: CONSTRAINT =
    (FORALL i: NAT | i < n)
        (EXISTS y: rm)
            Called_From(return_resource(i), y)

tx: CONSTRAINT =
    Called_From(the_tm.begin, the_ap)
    AND Called_From(the_tm.close, the_ap)
    AND Called_From(the_tm.commit, the_ap)
    AND Called_From(the_tm.information, the_ap)
    AND Called_From(the_tm.open, the_ap)
    AND Called_From(the_tm.rollback, the_ap)

xa: CONSTRAINT =
    (FORALL y: rm)
        [Called_From(the_tm.register, y)
         AND Called_From(the_tm.unregister, y)
         AND Called_From(y.close, the_tm)
         AND Called_From(y.commit, the_tm)
         AND Called_From(y.complete, the_tm)
         AND Called_From(y.end, the_tm)
         AND Called_From(y.forget, the_tm)
         AND Called_From(y.open, the_tm)
         AND Called_From(y.prepare, the_tm)
         AND Called_From(y.recover, the_tm)
         AND Called_From(y.rollback, the_tm)
         AND Called_From(y.start, the_tm)]
```

```
END x_open_proc_2
```

```

%%% Replace AR dataflow by a monitor, to make the communication asynchronous

x_open_proc_3: ARCHITECTURE [ -> ]

IMPORTING ALL FROM X_Open_style,
RPC_style

BEGIN

n: NAT      % Number of resource managers, a parameter in the specification

ar_requests, ar_resources: TYPE

q_type: {i: NAT | i < n} --> {t: TYPE | t < ar_requests}
r_type: {i: NAT | i < n} --> {t: TYPE | t < ar_resources}

COMPONENTS

ap: TYPE <= ARCHITECTURE [ -> ]

mon: TYPE <= ARCHITECTURE [ -> ]
      EXPORTING ALL
      BEGIN
          { put_requests(i): RPC [in: q_type(i) -> ] |(i: NAT) i < n }
          { get_requests(i): RPC [-> out: q_type(i)] |(i: NAT) i < n }
          { put_resources(i): RPC [in: r_type(i) -> ] |(i: NAT) i < n }
          { get_resources(i): RPC [-> out: r_type(i)] |(i: NAT) i < n }
      END mon

rm: TYPE <= { m: ARCHITECTURE [ -> ]
      EXPORTING ALL
      BEGIN
          close: XA_Close_Procedure
              [info: XA_Info, rmid: INT, flags: INT
               -> ret: INT]
          commit: XA_Commit_Procedure
              [id: X_Id, rmid: INT, flags: INT
               -> ret: INT]
          complete: XA_Complete_Procedure
              [hndl: INT, retval: INT,
               rmid: INT, flags: INT
               -> ret: INT]
          end: XA_End_Procedure
              [id: X_Id, rmid: INT, flags: INT
               -> ret: INT]
          forget: XA_Forget_Procedure
              [id: X_Id, rmid: INT, flags: INT
               -> ret: INT]
          open: XA_Open_Procedure
              [info: XA_Info, rmid: INT, flags: INT
               -> ret: INT]
          prepare: XA_Prepares_Procedure
              [id: X_Id, rmid: INT, flags: INT
               -> ret: INT]
          recover: XA_Recover_Procedure
              [ids: X_Ids, count: INT,
               rmid: INT, flags: INT
               -> ret: INT]
          rollback: XA_Rollback_Procedure
              [id: X_Id, rmid: INT, flags: INT
               -> ret: INT]
          start: XA_Start_Procedure
              [id: X_Id, rmid: INT, flags: INT
               -> ret: INT]
      END m
  END rm

```

```

    END m
    | qt < ar_requests AND rt < ar_resources }

rms: TYPE <= ARCHITECTURE [ -> ]

tm: TYPE <= ARCHITECTURE [ -> ]
    EXPORTING ALL
    BEGIN
        register: AX_Register_Procedure
            [id: X_Id, rmid: INT, flags: INT
             -> ret: INT]
        unregister: AX_Unregister_Procedure
            [rmid: INT, flags: INT
             -> ret: INT]
        begin: TX_Begin_Procedure [ -> ret: INT]
        close: TX_Close_Procedure [ -> ret: INT]
        commit: TX_Commit_Procedure [ -> ret: INT]
        information: TX_Info_Procedure [info: TX_Info -> ret: INT]
        open: TX_Open_Procedure [ -> ret: INT]
        rollback: TX_Rollback_Procedure [ -> ret: INT]
    END tm

the_ap: ap
the_rms: rms
the_tm: tm

CONFIGURATION

rms_contents: CONSTRAINT =
    (FORALL y: COMPONENT)
        [y PROPERLY_CONTAINED_IN the_rms => (EXISTS z: rm) y CONTAINED_IN z]
rm_location: CONSTRAINT =
    (FORALL y: rm) y CONTAINED_IN the_rms

ar_1: CONSTRAINT =
    (FORALL i: NAT | i < n)
        [Called_From(put_requests(i), the_ap)
         AND Called_From(get_resources(i), the_ap)]

ar_2: CONSTRAINT =
    (FORALL i: NAT | i < n)
        (EXISTS y: rm)
            [Called_From(put_resource(i), y)
             AND Called_From(get_requests(i), y)]

tx: CONSTRAINT =
    Called_From(the_tm.begin, the_ap)
    AND Called_From(the_tm.close, the_ap)
    AND Called_From(the_tm.commit, the_ap)
    AND Called_From(the_tm.information, the_ap)
    AND Called_From(the_tm.open, the_ap)
    AND Called_From(the_tm.rollback, the_ap)

xa: CONSTRAINT =
    (FORALL y: rm)
        [Called_From(the_tm.register, y)
         AND Called_From(the_tm.unregister, y)
         AND Called_From(y.close, the_tm)
         AND Called_From(y.commit, the_tm)
         AND Called_From(y.complete, the_tm)
         AND Called_From(y.end, the_tm)
         AND Called_From(y.forget, the_tm)
         AND Called_From(y.open, the_tm)
         AND Called_From(y.prepare, the_tm)]

```

```
AND Called_From(y.recover, the_tm)
AND Called_From(y.rollback, the_tm)
AND Called_From(y.start, the_tm)]  
END x_open_proc_3
```

```

%%% Pick a value for n, as a first step toward making things concrete

x_open_instance: ARCHITECTURE [ -> ]

IMPORTING ALL FROM X_Open_style,
RPC_style

BEGIN

ar_requests, ar_resources: TYPE

COMPONENTS

ap: TYPE <= ARCHITECTURE [ -> ]

rm: TYPE <= { m: ARCHITECTURE [ -> ]
              EXPORTING ALL
              BEGIN
                  access_function: RPC [in: qt -> out: rt]
                  close: XA_Close_Procedure
                      [info: XA_Info, rmid: INT, flags: INT
                       -> ret: INT]
                  commit: XA_Commit_Procedure
                      [id: X_Id, rmid: INT, flags: INT
                       -> ret: INT]
                  complete: XA_Complete_Procedure
                      [hdl: INT, retval: INT,
                       rmid: INT, flags: INT
                       -> ret: INT]
                  end: XA_End_Procedure
                      [id: X_Id, rmid: INT, flags: INT
                       -> ret: INT]
                  forget: XA_Forget_Procedure
                      [id: X_Id, rmid: INT, flags: INT
                       -> ret: INT]
                  open: XA_Open_Procedure
                      [info: XA_Info, rmid: INT, flags: INT
                       -> ret: INT]
                  prepare: XA_Prepares_Procedure
                      [id: X_Id, rmid: INT, flags: INT
                       -> ret: INT]
                  recover: XA_Recover_Procedure
                      [ids: X_Ids, count: INT,
                       rmid: INT, flags: INT
                       -> ret: INT]
                  rollback: XA_Rollback_Procedure
                      [id: X_Id, rmid: INT, flags: INT
                       -> ret: INT]
                  start: XA_Start_Procedure
                      [id: X_Id, rmid: INT, flags: INT
                       -> ret: INT]
              END m
              | qt < ar_requests AND rt < ar_resources }

rms: TYPE <= ARCHITECTURE [ -> ]

tm: TYPE <= ARCHITECTURE [ -> ]
      EXPORTING ALL
      BEGIN
          register: AX_Register_Procedure
              [id: X_Id, rmid: INT, flags: INT
               -> ret: INT]
          unregister: AX_Unregister_Procedure
              [rmid: INT, flags: INT
               -> ret: INT]

```

```

        -> ret: INT]
begin: TX_Begin_Procedure [ -> ret: INT]
close: TX_Close_Procedure [ -> ret: INT]
commit: TX_Commit_Procedure [ -> ret: INT]
information: TX_Info_Procedure [info: TX_Info -> ret: INT]
open: TX_Open_Procedure [ -> ret: INT]
rollback: TX_Rollback_Procedure [ -> ret: INT]
END tm

```

```

the_ap: ap
the_rms: rms
the_rm_1: rm
the_rm_2: rm
the_tm: tm

```

CONFIGURATION

```

rms_contents: CONSTRAINT =
  (FORALL y: COMPONENT)
    [y PROPERLY_CONTAINED_IN the_rms =>
     [y CONTAINED_IN the_rm_1 OR y CONTAINED_IN the_rm_2]]
rm_location: CONSTRAINT =
  the_rm_1 CONTAINED_IN the_rms AND the_rm_2 CONTAINED_IN the_rms
ar: CONSTRAINT =
  Called_From(the_rm_1.access_function, the_ap)
  AND Called_From(the_rm_2.access_function, the_ap)
tx: CONSTRAINT =
  Called_From(the_tm.begin, the_ap)
  AND Called_From(the_tm.close, the_ap)
  AND Called_From(the_tm.commit, the_ap)
  AND Called_From(the_tm.information, the_ap)
  AND Called_From(the_tm.open, the_ap)
  AND Called_From(the_tm.rollback, the_ap)
xa: CONSTRAINT =
  Called_From(the_tm.register, the_rm_1)
  AND Called_From(the_tm.unregister, the_rm_1)
  AND Called_From(the_rm_1.close, the_tm)
  AND Called_From(the_rm_1.commit, the_tm)
  AND Called_From(the_rm_1.complete, the_tm)
  AND Called_From(the_rm_1.end, the_tm)
  AND Called_From(the_rm_1.forget, the_tm)
  AND Called_From(the_rm_1.open, the_tm)
  AND Called_From(the_rm_1.prepare, the_tm)
  AND Called_From(the_rm_1.recover, the_tm)
  AND Called_From(the_rm_1.rollback, the_tm)
  AND Called_From(the_rm_1.start, the_tm)
  AND Called_From(the_tm.register, the_rm_2)
  AND Called_From(the_tm.unregister, the_rm_2)
  AND Called_From(the_rm_2.close, the_tm)
  AND Called_From(the_rm_2.commit, the_tm)
  AND Called_From(the_rm_2.complete, the_tm)
  AND Called_From(the_rm_2.end, the_tm)
  AND Called_From(the_rm_2.forget, the_tm)
  AND Called_From(the_rm_2.open, the_tm)
  AND Called_From(the_rm_2.prepare, the_tm)
  AND Called_From(the_rm_2.recover, the_tm)
  AND Called_From(the_rm_2.rollback, the_tm)
  AND Called_From(the_rm_2.start, the_tm)

```

```
END x_open_instance
```

```
%%% Break up AP into main process and an auxiliary process that will be
%%% co-located with the TM and the RMs. Step one is to refine example-6
%%% into something that looks like example-7, but with generic procedure
%%% calls in place of the RPCs. (So this could actually be above
%%% example-7 in the tree, rather than a separate branch.)
```

```
x_open_proc_4: ARCHITECTURE [ -> ]  
  IMPORTING ALL FROM X_Open_style  
  
BEGIN  
  
  n: NAT      % Number of resource managers, a parameter in the specification  
  ar_requests, ar_resources: TYPE  
  
COMPONENTS  
  
  ap: TYPE <= ARCHITECTURE [ -> ]  
  
  rm: TYPE <= { m: ARCHITECTURE [ -> ]  
    EXPORTING ALL  
    BEGIN  
      access_function: PROCEDURE [in: qt -> out: rt]  
      close: XA_Close_Procedure  
        [info: XA_Info, rmid: INT, flags: INT  
         -> ret: INT]  
      commit: XA_Commit_Procedure  
        [id: X_Id, rmid: INT, flags: INT  
         -> ret: INT]  
      complete: XA_Complete_Procedure  
        [hdl: INT, retval: INT,  
         rmid: INT, flags: INT  
         -> ret: INT]  
      end: XA_End_Procedure  
        [id: X_Id, rmid: INT, flags: INT  
         -> ret: INT]  
      forget: XA_Forget_Procedure  
        [id: X_Id, rmid: INT, flags: INT  
         -> ret: INT]  
      open: XA_Open_Procedure  
        [info: XA_Info, rmid: INT, flags: INT  
         -> ret: INT]  
      prepare: XA_Prepares_Procedure  
        [id: X_Id, rmid: INT, flags: INT  
         -> ret: INT]  
      recover: XA_Recover_Procedure  
        [ids: X_Ids, count: INT,  
         rmid: INT, flags: INT  
         -> ret: INT]  
      rollback: XA_Rollback_Procedure  
        [id: X_Id, rmid: INT, flags: INT  
         -> ret: INT]  
      start: XA_Start_Procedure  
        [id: X_Id, rmid: INT, flags: INT  
         -> ret: INT]  
    END m  
    | qt < ar_requests AND rt < ar_resources }  
  
  rms: TYPE <= ARCHITECTURE [ -> ]  
  tm: TYPE <= ARCHITECTURE [ -> ]  
    EXPORTING ALL  
    BEGIN
```

```

register: AX_Register_Procedure
    [id: X_Id, rmid: INT, flags: INT
     -> ret: INT]
unregister: AX_Unregister_Procedure
    [rmid: INT, flags: INT
     -> ret: INT]
begin: TX_Begin_Procedure [ -> ret: INT]
close: TX_Close_Procedure [ -> ret: INT]
commit: TX_Commit_Procedure [ -> ret: INT]
information: TX_Info_Procedure [info: TX_Info -> ret: INT]
open: TX_Open_Procedure [ -> ret: INT]
rollback: TX_Rollback_Procedure [ -> ret: INT]
END tm

the_ap: ap
the_rms: rms
the_tm: tm

CONFIGURATION

rms_contents: CONSTRAINT =
    (FORALL y: COMPONENT)
        [y PROPERLY_CONTAINED_IN the_rms => (EXISTS z: rm) y CONTAINED_IN z]
rm_location: CONSTRAINT =
    (FORALL y: rm) y CONTAINED_IN the_rms

ar: CONSTRAINT =
    (FORALL y: rm) Called_From(y.access_function, the_ap)

tx: CONSTRAINT =
    Called_From(the_tm.begin, the_ap)
    AND Called_From(the_tm.close, the_ap)
    AND Called_From(the_tm.commit, the_ap)
    AND Called_From(the_tm.information, the_ap)
    AND Called_From(the_tm.open, the_ap)
    AND Called_From(the_tm.rollback, the_ap)

xa: CONSTRAINT =
    (FORALL y: rm)
        [Called_From(the_tm.register, y)
         AND Called_From(the_tm.unregister, y)
         AND Called_From(y.close, the_tm)
         AND Called_From(y.commit, the_tm)
         AND Called_From(y.complete, the_tm)
         AND Called_From(y.end, the_tm)
         AND Called_From(y.forget, the_tm)
         AND Called_From(y.open, the_tm)
         AND Called_From(y.prepare, the_tm)
         AND Called_From(y.recover, the_tm)
         AND Called_From(y.rollback, the_tm)
         AND Called_From(y.start, the_tm)]]

END x_open_proc_4

```

%%% Break up AP into main process and an auxiliary process that will be
%%% co-located with the TM and the RMs. This is an intermediate step in
%%% which the various boxes representing the auxiliary interface processes
%%% for the AP and the TM are each combined into a single box.

```
x_open_ap_decomposition: ARCHITECTURE [ -> ]  
  
IMPORTING ALL FROM X_Open_style, RPC_style  
  
BEGIN  
  
n: NAT      % Number of resource managers, a parameter in the specification  
ar_requests, ar_resources: TYPE  
  
resource_id: TYPE = { i: NAT | i < n }
```

COMPONENTS

%% The next two type definitions could simply be declared within the ap
%% declaration, since the number is fixed, but might as well do it like rm.

```
ap_main: TYPE <= ARCHITECTURE [ -> ]  
  
ap_aux:  TYPE <= ARCHITECTURE [ -> ]  
          EXPORTING ALL  
          BEGIN  
              parameterized_access_function:  
                  RPC [r_id: resource_id, in: ar_requests -> out: ar_resources]  
                  begin: RPC [ -> ret: INT]  
                  close: RPC [ -> ret: INT]  
                  commit: RPC [ -> ret: INT]  
                  information: RPC [info: TX_Info -> ret: INT]  
                  open: RPC [ -> ret: INT]  
                  rollback: RPC [ -> ret: INT]  
          END ap_aux  
  
ap: TYPE <= ARCHITECTURE [ -> ]  
  
rm: TYPE <= { m: ARCHITECTURE [ -> ]  
          EXPORTING ALL  
          BEGIN  
              access_function: PROCEDURE [in: qt -> out: rt]  
              close: XA_Close_Procedure  
                  [info: XA_Info, rmid: INT, flags: INT  
                   -> ret: INT]  
              commit: XA_Commit_Procedure  
                  [id: X_Id, rmid: INT, flags: INT  
                   -> ret: INT]  
              complete: XA_Complete_Procedure  
                  [hdl1: INT, retval: INT,  
                   rmid: INT, flags: INT  
                   -> ret: INT]  
              end: XA_End_Procedure  
                  [id: X_Id, rmid: INT, flags: INT  
                   -> ret: INT]  
              forget: XA_Forget_Procedure  
                  [id: X_Id, rmid: INT, flags: INT  
                   -> ret: INT]  
              open: XA_Open_Procedure  
                  [info: XA_Info, rmid: INT, flags: INT  
                   -> ret: INT]  
              prepare: XA_Prepares_Procedure  
                  [id: X_Id, rmid: INT, flags: INT]
```

```

        -> ret: INT]
recover: XA_Recover_Procedure
    [ids: X_Ids, count: INT,
     rmid: INT, flags: INT
     -> ret: INT]
rollback: XA_Rollback_Procedure
    [id: X_Id, rmid: INT, flags: INT
     -> ret: INT]
start: XA_Start_Procedure
    [id: X_Id, rmid: INT, flags: INT
     -> ret: INT]
END m
| qt < ar_requests AND rt < ar_resources }

rms: TYPE <= ARCHITECTURE [ -> ]

tm_main: TYPE <= ARCHITECTURE [ -> ]
EXPORTING ALL
BEGIN
    register: RPC
        [id: X_Id, rmid: INT, flags: INT
         -> ret: INT]
    unregister: RPC
        [rmid: INT, flags: INT
         -> ret: INT]
    begin: TX_Begin_Procedure [ -> ret: INT]
    close: TX_Close_Procedure [ -> ret: INT]
    commit: TX_Commit_Procedure [ -> ret: INT]
    information: TX_Info_Procedure [info: TX_Info -> ret: INT]
    open: TX_Open_Procedure [ -> ret: INT]
    rollback: TX_Rollback_Procedure [ -> ret: INT]
END tm_main

tm_aux: TYPE <= ARCHITECTURE [ -> ]
EXPORTING ALL
BEGIN
    register: AX_Register_Procedure
        [id: X_Id, rmid: INT, flags: INT
         -> ret: INT]
    unregister: AX_Unregister_Procedure
        [rmid: INT, flags: INT
         -> ret: INT]
    close: RPC [info: XA_Info, rmid: INT, flags: INT
               -> ret: INT]
    commit: RPC [id: X_Id, rmid: INT, flags: INT
                 -> ret: INT]
    complete: RPC [hdl: INT, retval: INT,
                  rmid: INT, flags: INT
                  -> ret: INT]
    end: RPC [id: X_Id, rmid: INT, flags: INT
              -> ret: INT]
    forget: RPC [id: X_Id, rmid: INT, flags: INT
                 -> ret: INT]
    open: RPC [info: XA_Info, rmid: INT, flags: INT
               -> ret: INT]
    prepare: RPC [id: X_Id, rmid: INT, flags: INT
                  -> ret: INT]
    recover: RPC [ids: X_Ids, count: INT,
                  rmid: INT, flags: INT
                  -> ret: INT]
    rollback: RPC [id: X_Id, rmid: INT, flags: INT
                  -> ret: INT]
    start: RPC [id: X_Id, rmid: INT, flags: INT
                 -> ret: INT]

```

```

        END tm_aux

tm: TYPE <= ARCHITECTURE [ -> ]

the_ap: ap
the_ap_main: ap_main
the_ap_aux: ap_aux
the_rms: rms
the_tm: tm
the_tm_main: tm_main
the_tm_aux: tm_aux

CONFIGURATION

ap_contents: CONSTRAINT =
  (FORALL y: COMPONENT)
    [y PROPERLY_CONTAINED_IN the_ap
     => y CONTAINED_IN the_ap_main OR y CONTAINED_IN the_ap_aux]
ap_main_location: CONSTRAINT =
  the_ap_main CONTAINED_IN the_ap
ap_aux_location: CONSTRAINT =
  the_ap_aux CONTAINED_IN the_ap

tm_contents: CONSTRAINT =
  (FORALL y: COMPONENT)
    [y PROPERLY_CONTAINED_IN the_tm
     => y CONTAINED_IN the_tm_main OR y CONTAINED_IN the_tm_aux]
tm_main_location: CONSTRAINT =
  the_tm_main CONTAINED_IN the_tm
tm_aux_location: CONSTRAINT =
  the_tm_aux CONTAINED_IN the_tm

rms_contents: CONSTRAINT =
  (FORALL y: COMPONENT)
    [y PROPERLY_CONTAINED_IN the_rms => (EXISTS z: rm) y CONTAINED_IN z]
rm_location: CONSTRAINT =
  (FORALL y: rm) y CONTAINED_IN the_rms

ar: CONSTRAINT =
  (FORALL y: rm) Called_From(y.access_function, the_ap_aux)

tx: CONSTRAINT =
  Called_From(the_tm_main.begin, the_ap_aux)
  AND Called_From(the_tm_main.close, the_ap_aux)
  AND Called_From(the_tm_main.commit, the_ap_aux)
  AND Called_From(the_tm_main.information, the_ap_aux)
  AND Called_From(the_tm_main.open, the_ap_aux)
  AND Called_From(the_tm_main.rollback, the_ap_aux)

xa: CONSTRAINT =
  (FORALL y: rm)
    [Called_From(the_tm_aux.register, y)
     AND Called_From(the_tm_aux.unregister, y)
     AND Called_From(y.close, the_tm_aux)
     AND Called_From(y.commit, the_tm_aux)
     AND Called_From(y.complete, the_tm_aux)
     AND Called_From(y.end, the_tm_aux)
     AND Called_From(y.forget, the_tm_aux)
     AND Called_From(y.open, the_tm_aux)
     AND Called_From(y.prepare, the_tm_aux)
     AND Called_From(y.recover, the_tm_aux)
     AND Called_From(y.rollback, the_tm_aux)
     AND Called_From(y.start, the_tm_aux)]

```

```
intra_ap: CONSTRAINT =
    Called_From(parameterized_access_function.the_ap_aux, the_ap)
    AND Called_From(the_ap_aux.begin, the_ap)
    AND Called_From(the_ap_aux.close, the_ap)
    AND Called_From(the_ap_aux.commit, the_ap)
    AND Called_From(the_ap_aux.information, the_ap)
    AND Called_From(the_ap_aux.open, the_ap)
    AND Called_From(the_ap_aux.rollback, the_ap)

intra_tm: CONSTRAINT =
    Called_From(the_tm_main.register, the_tm_aux)
    AND Called_From(the_tm_main.unregister, the_tm_aux)
    AND Called_From(the_tm_aux.close, the_tm_main)
    AND Called_From(the_tm_aux.commit, the_tm_main)
    AND Called_From(the_tm_aux.complete, the_tm_main)
    AND Called_From(the_tm_aux.end, the_tm_main)
    AND Called_From(the_tm_aux.forget, the_tm_main)
    AND Called_From(the_tm_aux.open, the_tm_main)
    AND Called_From(the_tm_aux.prepare, the_tm_main)
    AND Called_From(the_tm_aux.recover, the_tm_main)
    AND Called_From(the_tm_aux.rollback, the_tm_main)
    AND Called_From(the_tm_aux.start, the_tm_main)

END x_open_ap_decomposition
```

```

%%% Break out interfaces of ap_ar_aux and tm_aux to handle distributed RMs.

x_open_ap_aux_decomposition: ARCHITECTURE [ -> ]

    IMPORTING ALL FROM X_Open_style, RPC_style

BEGIN

    n: NAT      % Number of resource managers, a parameter in the specification
    ar_requests, ar_resources: TYPE

COMPONENTS

    ap_main: TYPE <= ARCHITECTURE [ -> ]

    ap_tx_aux:  TYPE <= ARCHITECTURE [ -> ]
        EXPORTING ALL
        BEGIN
            begin: RPC [ -> ret: INT]
            close: RPC [ -> ret: INT]
            commit: RPC [ -> ret: INT]
            information: RPC [info: TX_Info -> ret: INT]
            open: RPC [ -> ret: INT]
            rollback: RPC [ -> ret: INT]
        END ap_tx_aux

    ap_ar_aux: TYPE <= {m: ARCHITECTURE [ -> ]
        EXPORTING ALL
        BEGIN
            access_function:
                RPC [in: qt -> out: rt]
        END m
        | qt < ar_requests AND rt < ar_resources }

ap: TYPE <= ARCHITECTURE [ -> ]

rm: TYPE <= { m: ARCHITECTURE [ -> ]
    EXPORTING ALL
    BEGIN
        access_function: PROCEDURE [in: qt -> out: rt]
        close: XA_Close_Procedure
            [info: XA_Info, rmid: INT, flags: INT
             -> ret: INT]
        commit: XA_Commit_Procedure
            [id: X_Id, rmid: INT, flags: INT
             -> ret: INT]
        complete: XA_Complete_Procedure
            [hdl: INT, retval: INT,
             rmid: INT, flags: INT
             -> ret: INT]
        end: XA_End_Procedure
            [id: X_Id, rmid: INT, flags: INT
             -> ret: INT]
        forget: XA_Forget_Procedure
            [id: X_Id, rmid: INT, flags: INT
             -> ret: INT]
        open: XA_Open_Procedure
            [info: XA_Info, rmid: INT, flags: INT
             -> ret: INT]
        prepare: XA_Prepares_Procedure
            [id: X_Id, rmid: INT, flags: INT
             -> ret: INT]
        recover: XA_Recover_Procedure
    
```

```

        [ids: X_Ids, count: INT,
         rmid: INT, flags: INT
         -> ret: INT]
    rollback: XA_Rollback_Procedure
        [id: X_Id, rmid: INT, flags: INT
         -> ret: INT]
    start: XA_Start_Procedure
        [id: X_Id, rmid: INT, flags: INT
         -> ret: INT]
END m
| qt < ar_requests AND rt < ar_resources }

rms: TYPE <= ARCHITECTURE [ -> ]

tm_main: TYPE <= ARCHITECTURE [ -> ]
EXPORTING ALL
BEGIN
    register: RPC
        [id: X_Id, rmid: INT, flags: INT
         -> ret: INT]
    unregister: RPC
        [rmid: INT, flags: INT
         -> ret: INT]
    begin: TX_Begin_Procedure [ -> ret: INT]
    close: TX_Close_Procedure [ -> ret: INT]
    commit: TX_Commit_Procedure [ -> ret: INT]
    information: TX_Info_Procedure [info: TX_Info -> ret: INT]
    open: TX_Open_Procedure [ -> ret: INT]
    rollback: TX_Rollback_Procedure [ -> ret: INT]
END tm_main

tm_aux: TYPE <= ARCHITECTURE [ -> ]
EXPORTING ALL
BEGIN
    register: AX_Register_Procedure
        [id: X_Id, rmid: INT, flags: INT
         -> ret: INT]
    unregister: AX_Unregister_Procedure
        [rmid: INT, flags: INT
         -> ret: INT]
    close: RPC [info: XA_Info, rmid: INT, flags: INT
               -> ret: INT]
    commit: RPC [id: X_Id, rmid: INT, flags: INT
               -> ret: INT]
    complete: RPC [hdl: INT, retval: INT,
                  rmid: INT, flags: INT
                  -> ret: INT]
    end: RPC [id: X_Id, rmid: INT, flags: INT
              -> ret: INT]
    forget: RPC [id: X_Id, rmid: INT, flags: INT
                 -> ret: INT]
    open: RPC [info: XA_Info, rmid: INT, flags: INT
               -> ret: INT]
    prepare: RPC [id: X_Id, rmid: INT, flags: INT
                  -> ret: INT]
    recover: RPC [ids: X_Ids, count: INT,
                  rmid: INT, flags: INT
                  -> ret: INT]
    rollback: RPC [id: X_Id, rmid: INT, flags: INT
                  -> ret: INT]
    start: RPC [id: X_Id, rmid: INT, flags: INT
                 -> ret: INT]
END tm_aux

```

```

tm: TYPE <= ARCHITECTURE [ -> ]

the_ap: ap
the_ap_main: ap_main
the_ap_tx_aux: ap_tx_aux
the_rms: rms
the_tm: tm
the_tm_main: tm_main

CONFIGURATION

ap_contents: CONSTRAINT =
  (FORALL y: COMPONENT)
    [y PROPERLY_CONTAINED_IN the_ap
      => y CONTAINED_IN the_ap_main
      OR y CONTAINED_IN the_ap_tx_aux
      OR (EXISTS w: ap_ar_aux) y CONTAINED_IN w]
ap_main_location: CONSTRAINT =
  the_ap_main CONTAINED_IN the_ap
tx_aux_location: CONSTRAINT =
  the_ap_tx_aux CONTAINED_IN the_ap
ra_aux_location: CONSTRAINT =
  (FORALL w: ap_ar_aux) w CONTAINED_IN the_ap

tm_contents: CONSTRAINT =
  (FORALL y: COMPONENT)
    [y PROPERLY_CONTAINED_IN the_tm
      => y CONTAINED_IN the_tm_main
      OR (EXISTS w: tm_aux) y CONTAINED_IN w]
tm_main_location: CONSTRAINT =
  the_tm_main CONTAINED_IN the_tm
tm_aux_location: CONSTRAINT =
  (FORALL w: tm_aux) w CONTAINED_IN the_tm

rms_contents: CONSTRAINT =
  (FORALL y: COMPONENT)
    [y PROPERLY_CONTAINED_IN the_rms => (EXISTS z: rm) y CONTAINED_IN z]
rm_location: CONSTRAINT =
  (FORALL y: rm) y CONTAINED_IN the_rms

ar_1: CONSTRAINT =
  (FORALL y: rm) (EXISTS w: ap_ar_aux) Called_From(y.access_function, w)
ar_2: CONSTRAINT =
  (FORALL w: ap_ar_aux) (EXISTS y: rm) Called_From(y.access_function, w)

tx: CONSTRAINT =
  Called_From(the_tm_main.begin, the_ap_tx_aux)
  AND Called_From(the_tm_main.close, the_ap_tx_aux)
  AND Called_From(the_tm_main.commit, the_ap_tx_aux)
  AND Called_From(the_tm_main.information, the_ap_tx_aux)
  AND Called_From(the_tm_main.open, the_ap_tx_aux)
  AND Called_From(the_tm_main.rollback, the_ap_tx_aux)

xa_1: CONSTRAINT =
  (FORALL y: rm) (EXISTS w: tm_aux)
  [Called_From(w.register, y)
    AND Called_From(w.unregister, y)
    AND Called_From(y.close, w)
    AND Called_From(y.commit, w)
    AND Called_From(y.complete, w)
    AND Called_From(y.end, w)
    AND Called_From(y.forget, w)]

```

```

        AND Called_From(y.open, w)
        AND Called_From(y.prepare, w)
        AND Called_From(y.recover, w)
        AND Called_From(y.rollback, w)
        AND Called_From(y.start, w)]
```

xa_2: CONSTRAINT =
 (FORALL w: tm_aux) (EXISTS y: rm)
 [Called_From(w.register, y)
 AND Called_From(w.unregister, y)
 AND Called_From(y.close, w)
 AND Called_From(y.commit, w)
 AND Called_From(y.complete, w)
 AND Called_From(y.end, w)
 AND Called_From(y.forget, w)
 AND Called_From(y.open, w)
 AND Called_From(y.prepare, w)
 AND Called_From(y.recover, w)
 AND Called_From(y.rollback, w)
 AND Called_From(y.start, w)]

intra_ap_1: CONSTRAINT =
 (FORALL w: ap_ar_aux) Called_From(w.access_function, the_ap_main)

intra_ap_2: CONSTRAINT =
 Called_From(the_ap_tx_aux.begin, the_ap)
 AND Called_From(the_ap_tx_aux.close, the_ap)
 AND Called_From(the_ap_tx_aux.commit, the_ap)
 AND Called_From(the_ap_tx_aux.information, the_ap)
 AND Called_From(the_ap_tx_aux.open, the_ap)
 AND Called_From(the_ap_tx_aux.rollback, the_ap)

intra_tm: CONSTRAINT =
 (FORALL w: tm_aux)
 [Called_From(the_tm_main.register, w)
 AND Called_From(the_tm_main.unregister, w)
 AND Called_From(w.close, the_tm_main)
 AND Called_From(w.commit, the_tm_main)
 AND Called_From(w.complete, the_tm_main)
 AND Called_From(w.end, the_tm_main)
 AND Called_From(w.forget, the_tm_main)
 AND Called_From(w.open, the_tm_main)
 AND Called_From(w.prepare, the_tm_main)
 AND Called_From(w.recover, the_tm_main)
 AND Called_From(w.rollback, the_tm_main)
 AND Called_From(w.start, the_tm_main)]

END x_open_ap_aux_decomposition

```

%%% Now we shoot for the dual to example-13, breaking out the auxiliary
%%% processes on the RM side rather than the AP side.

x_open_manager_decomposition: ARCHITECTURE [ -> ]

IMPORTING ALL FROM X_Open_style

BEGIN

n: NAT      % Number of resource managers, a parameter in the specification

ar_requests, ar_resources: TYPE

COMPONENTS

ap: TYPE <= ARCHITECTURE [ -> ]

rm_ar_aux: TYPE <= { m: ARCHITECTURE [ -> ]
                      EXPORTING ALL
                      BEGIN
                        access_function: PROCEDURE [in: qt -> out: rt]
                      END m
                      | qt < ar_requests AND rt < ar_resources }

rm_xa_aux: TYPE <= ARCHITECTURE [ -> ]
                      EXPORTING ALL
                      BEGIN
                        close: XA_Close_Procedure
                            [info: XA_Info, rmid: INT, flags: INT
                             -> ret: INT]
                        commit: XA_Commit_Procedure
                            [id: X_Id, rmid: INT, flags: INT
                             -> ret: INT]
                        complete: XA_Complete_Procedure
                            [hdl: INT, retval: INT,
                             rmid: INT, flags: INT
                             -> ret: INT]
                        end: XA_End_Procedure
                            [id: X_Id, rmid: INT, flags: INT
                             -> ret: INT]
                        forget: XA_Forget_Procedure
                            [id: X_Id, rmid: INT, flags: INT
                             -> ret: INT]
                        open: XA_Open_Procedure
                            [info: XA_Info, rmid: INT, flags: INT
                             -> ret: INT]
                        prepare: XA_Prepares_Procedure
                            [id: X_Id, rmid: INT, flags: INT
                             -> ret: INT]
                        recover: XA_Recover_Procedure
                            [ids: X_Ids, count: INT,
                             rmid: INT, flags: INT
                             -> ret: INT]
                        rollback: XA_Rollback_Procedure
                            [id: X_Id, rmid: INT, flags: INT
                             -> ret: INT]
                        start: XA_Start_Procedure
                            [id: X_Id, rmid: INT, flags: INT
                             -> ret: INT]
                      END rm_xa_aux

rm_main: TYPE <= { m: ARCHITECTURE [ -> ]
                      EXPORTING ALL
                      BEGIN

```

```

access_function: RPC [in: qt -> out: rt]
close: RPC
    [info: XA_Info, rmid: INT, flags: INT
     -> ret: INT]
commit: RPC
    [id: X_Id, rmid: INT, flags: INT
     -> ret: INT]
complete: RPC
    [hdl: INT, retval: INT,
     rmid: INT, flags: INT
     -> ret: INT]
end: RPC
    [id: X_Id, rmid: INT, flags: INT
     -> ret: INT]
forget: RPC
    [id: X_Id, rmid: INT, flags: INT
     -> ret: INT]
open: RPC
    [info: XA_Info, rmid: INT, flags: INT
     -> ret: INT]
prepare: RPC
    [id: X_Id, rmid: INT, flags: INT
     -> ret: INT]
recover: RPC
    [ids: X_Ids, count: INT,
     rmid: INT, flags: INT
     -> ret: INT]
rollback: RPC
    [id: X_Id, rmid: INT, flags: INT
     -> ret: INT]
start: RPC
    [id: X_Id, rmid: INT, flags: INT
     -> ret: INT]
END m
| qt < ar_requests AND rt < ar_resources }

rms: TYPE <= ARCHITECTURE [ -> ]

tm_aux: TYPE <= ARCHITECTURE [ -> ]
    EXPORTING ALL
    BEGIN
        begin: TX_Begin_Procedure [ -> ret: INT]
        close: TX_Close_Procedure [ -> ret: INT]
        commit: TX_Commit_Procedure [ -> ret: INT]
        information: TX_Info_Procedure [info: TX_Info -> ret: INT]
        open: TX_Open_Procedure [ -> ret: INT]
        rollback: TX_Rollback_Procedure [ -> ret: INT]
    END tm_aux

tm_main: TYPE <= ARCHITECTURE [ -> ]
    EXPORTING ALL
    BEGIN
        register: AX_Register_Procedure
            [id: X_Id, rmid: INT, flags: INT
             -> ret: INT]
        unregister: AX_Unregister_Procedure
            [rmid: INT, flags: INT
             -> ret: INT]
        begin: RPC [ -> ret: INT]
        close: RPC [ -> ret: INT]
        commit: RPC [ -> ret: INT]
        information: RPC [info: TX_Info -> ret: INT]
        open: RPC [ -> ret: INT]
        rollback: RPC [ -> ret: INT]

```

```

END tm_main

the_ap: ap
the_rms: rms
the_tm: tm
the_tm_main: tm_main
the_tm_aux: tm_aux

CONFIGURATION

rms_contents: CONSTRAINT =
  (FORALL y: COMPONENT)
    [y PROPERLY_CONTAINED_IN the_rms
     => (EXISTS z: rm_main) y CONTAINED_IN z
        OR (EXISTS z: rm_ar_aux) y CONTAINED_IN z
        OR (EXISTS z: rm_xa_aux) y CONTAINED_IN z]
rm_main_location: CONSTRAINT =
  (FORALL y: rm_main) y CONTAINED_IN the_rms
rm_ar_aux_location: CONSTRAINT =
  (FORALL y: rm_ar_aux) y CONTAINED_IN the_rms
rm_xa_aux_location: CONSTRAINT =
  (FORALL y: rm_xa_aux) y CONTAINED_IN the_rms

ar: CONSTRAINT =
  (FORALL u: rm_ar_aux) Called_From(u.access_function, the_ap)

tx: CONSTRAINT =
  Called_From(the_tm_aux.begin, the_ap)
  AND Called_From(the_tm_aux.close, the_ap)
  AND Called_From(the_tm_aux.commit, the_ap)
  AND Called_From(the_tm_aux.information, the_ap)
  AND Called_From(the_tm_aux.open, the_ap)
  AND Called_From(the_tm_aux.rollback, the_ap)

xa: CONSTRAINT =
  (FORALL v: rm_xa_aux)
    [Called_From(the_tm.register, v)
     AND Called_From(the_tm.unregister, v)
     AND Called_From(v.close, the_tm)
     AND Called_From(v.commit, the_tm)
     AND Called_From(v.complete, the_tm)
     AND Called_From(v.end, the_tm)
     AND Called_From(v.forget, the_tm)
     AND Called_From(v.open, the_tm)
     AND Called_From(v.prepare, the_tm)
     AND Called_From(v.recover, the_tm)
     AND Called_From(v.rollback, the_tm)
     AND Called_From(v.start, the_tm)]]

intra_tm: CONSTRAINT =
  Called_From(the_tm_main.begin, the_tm_aux)
  AND Called_From(the_tm_main.close, the_tm_aux)
  AND Called_From(the_tm_main.commit, the_tm_aux)
  AND Called_From(the_tm_main.information, the_tm_aux)
  AND Called_From(the_tm_main.open, the_tm_aux)
  AND Called_From(the_tm_main.rollback, the_tm_aux)

intra_rm_1: CONSTRAINT =
  (FORALL y: rm_main) (EXISTS u: rm_ar_aux)
    Called_From(y.access_function, u)

intra_rm_2: CONSTRAINT =
  (FORALL u: rm_ar_aux) (EXISTS y: rm_main)

```

```

    Called_From(y.access_function, u)

intra_rm_3: CONSTRAINT =
  (FORALL y: rm_main) (EXISTS v: rm_xa_aux)
    [Called_From(v.register, y)
     AND Called_From(v.unregister, y)
     AND Called_From(y.close, v)
     AND Called_From(y.commit, v)
     AND Called_From(y.complete, v)
     AND Called_From(y.end, v)
     AND Called_From(y.forget, v)
     AND Called_From(y.open, v)
     AND Called_From(y.prepare, v)
     AND Called_From(y.recover, v)
     AND Called_From(y.rollback, v)
     AND Called_From(y.start, v)]

intra_rm_4: CONSTRAINT =
  (FORALL v: rm_xa_aux) (EXISTS y: rm_main)
    [Called_From(v.register, y)
     AND Called_From(v.unregister, y)
     AND Called_From(y.close, v)
     AND Called_From(y.commit, v)
     AND Called_From(y.complete, v)
     AND Called_From(y.end, v)
     AND Called_From(y.forget, v)
     AND Called_From(y.open, v)
     AND Called_From(y.prepare, v)
     AND Called_From(y.recover, v)
     AND Called_From(y.rollback, v)
     AND Called_From(y.start, v)]

END x_open_manager_decomposition

```

%%% Similar to example-14, but collapses auxiliary processes in RM. Dual
%%% to example-12 in much the way example-14 is dual to example-13.

```
x_open_manager_alt_decomposition: ARCHITECTURE [ -> ]  
  
    IMPORTING ALL FROM X_Open_style  
  
BEGIN  
  
    n: NAT      % Number of resource managers, a parameter in the specification  
  
    ar_requests, ar_resources: TYPE  
  
    resource_id: TYPE = { i: NAT | i < n }  
  
COMPONENTS  
  
    ap: TYPE <= ARCHITECTURE [ -> ]  
  
    rm_ar_aux: TYPE <= ARCHITECTURE [ -> ]  
        EXPORTING ALL  
        BEGIN  
            parameterized_access_function:  
                PROCEDURE [r_id: resource_id, in: qt -> out: rt]  
        END rm_ar_aux  
  
%% Note that the following already contain the necessary RM id args.  
rm_xa_aux: TYPE <= ARCHITECTURE [ -> ]  
    EXPORTING ALL  
    BEGIN  
        close: XA_Close_Procedure  
            [info: XA_Info, rmid: INT, flags: INT  
             -> ret: INT]  
        commit: XA_Commit_Procedure  
            [id: X_Id, rmid: INT, flags: INT  
             -> ret: INT]  
        complete: XA_Complete_Procedure  
            [hndl: INT, retval: INT,  
             rmid: INT, flags: INT  
             -> ret: INT]  
        end: XA_End_Procedure  
            [id: X_Id, rmid: INT, flags: INT  
             -> ret: INT]  
        forget: XA_Forget_Procedure  
            [id: X_Id, rmid: INT, flags: INT  
             -> ret: INT]  
        open: XA_Open_Procedure  
            [info: XA_Info, rmid: INT, flags: INT  
             -> ret: INT]  
        prepare: XA_Prepares_Procedure  
            [id: X_Id, rmid: INT, flags: INT  
             -> ret: INT]  
        recover: XA_Recover_Procedure  
            [ids: X_Ids, count: INT,  
             rmid: INT, flags: INT  
             -> ret: INT]  
        rollback: XA_Rollback_Procedure  
            [id: X_Id, rmid: INT, flags: INT  
             -> ret: INT]  
        start: XA_Start_Procedure  
            [id: X_Id, rmid: INT, flags: INT  
             -> ret: INT]  
    END rm_xa_aux
```

```

rm_main: TYPE <= { m: ARCHITECTURE [ -> ]
                    EXPORTING ALL
                    BEGIN
                        access_function: RPC [in: qt -> out: rt]
                        close: RPC
                            [info: XA_Info, rmid: INT, flags: INT
                             -> ret: INT]
                        commit: RPC
                            [id: X_Id, rmid: INT, flags: INT
                             -> ret: INT]
                        complete: RPC
                            [hdl: INT, retval: INT,
                             rmid: INT, flags: INT
                             -> ret: INT]
                        end: RPC
                            [id: X_Id, rmid: INT, flags: INT
                             -> ret: INT]
                        forget: RPC
                            [id: X_Id, rmid: INT, flags: INT
                             -> ret: INT]
                        open: RPC
                            [info: XA_Info, rmid: INT, flags: INT
                             -> ret: INT]
                        prepare: RPC
                            [id: X_Id, rmid: INT, flags: INT
                             -> ret: INT]
                        recover: RPC
                            [ids: X_Ids, count: INT,
                             rmid: INT, flags: INT
                             -> ret: INT]
                        rollback: RPC
                            [id: X_Id, rmid: INT, flags: INT
                             -> ret: INT]
                        start: RPC
                            [id: X_Id, rmid: INT, flags: INT
                             -> ret: INT]
                    END m
                    | qt < ar_requests AND rt < ar_resources }

rms: TYPE <= ARCHITECTURE [ -> ]

tm_aux: TYPE <= ARCHITECTURE [ -> ]
                    EXPORTING ALL
                    BEGIN
                        begin: TX_Begin_Procedure [ -> ret: INT]
                        close: TX_Close_Procedure [ -> ret: INT]
                        commit: TX_Commit_Procedure [ -> ret: INT]
                        information: TX_Info_Procedure [info: TX_Info -> ret: INT]
                        open: TX_Open_Procedure [ -> ret: INT]
                        rollback: TX_Rollback_Procedure [ -> ret: INT]
                    END tm_aux

tm_main: TYPE <= ARCHITECTURE [ -> ]
                    EXPORTING ALL
                    BEGIN
                        register: AX_Register_Procedure
                            [id: X_Id, rmid: INT, flags: INT
                             -> ret: INT]
                        unregister: AX_Unregister_Procedure
                            [rmid: INT, flags: INT
                             -> ret: INT]
                        begin: RPC [ -> ret: INT]
                        close: RPC [ -> ret: INT]
                        commit: RPC [ -> ret: INT]

```

```

information: RPC [info: TX_Info -> ret: INT]
open: RPC [ -> ret: INT]
rollback: RPC [ -> ret: INT]
END tm_main

the_ap: ap
the_rms: rms
the_rm_ar_aux: rm_ar_aux
the_rm_xa_aux: rm_xa_aux
the_tm: tm
the_tm_main: tm_main
the_tm_aux: tm_aux

CONFIGURATION

rms_contents: CONSTRAINT =
(FORALL y: COMPONENT)
[y PROPERLY_CONTAINED_IN the_rms
=> (EXISTS z: rm_main) y CONTAINED_IN z
OR y CONTAINED_IN the_rm_ar_aux
OR y CONTAINED_IN the_rm_xa_aux]

rm_main_location: CONSTRAINT =
(FORALL y: rm_main) y CONTAINED_IN the_rms

rm_ar_aux_location: CONSTRAINT =
the_rm_ar_aux CONTAINED_IN the_rms

rm_xa_aux_location: CONSTRAINT =
the_rm_xa_aux CONTAINED_IN the_rms

ar: CONSTRAINT =
Called_From(the_rm_ar_aux.access_function, the_ap)

tx: CONSTRAINT =
Called_From(the_tm_aux.begin, the_ap)
AND Called_From(the_tm_aux.close, the_ap)
AND Called_From(the_tm_aux.commit, the_ap)
AND Called_From(the_tm_aux.information, the_ap)
AND Called_From(the_tm_aux.open, the_ap)
AND Called_From(the_tm_aux.rollback, the_ap)

xa: CONSTRAINT =
Called_From(the_tm.register, the_rm_xa_aux)
AND Called_From(the_tm.unregister, the_rm_xa_aux)
AND Called_From(the_rm_xa_aux.close, the_tm)
AND Called_From(the_rm_xa_aux.commit, the_tm)
AND Called_From(the_rm_xa_aux.complete, the_tm)
AND Called_From(the_rm_xa_aux.end, the_tm)
AND Called_From(the_rm_xa_aux.forget, the_tm)
AND Called_From(the_rm_xa_aux.open, the_tm)
AND Called_From(the_rm_xa_aux.prepare, the_tm)
AND Called_From(the_rm_xa_aux.recover, the_tm)
AND Called_From(the_rm_xa_aux.rollback, the_tm)
AND Called_From(the_rm_xa_aux.start, the_tm)

intra_tm: CONSTRAINT =
Called_From(the_tm_main.begin, the_tm_aux)
AND Called_From(the_tm_main.close, the_tm_aux)
AND Called_From(the_tm_main.commit, the_tm_aux)
AND Called_From(the_tm_main.information, the_tm_aux)
AND Called_From(the_tm_main.open, the_tm_aux)
AND Called_From(the_tm_main.rollback, the_tm_aux)

intra_rm_1: CONSTRAINT =
(FORALL y: rm_main)

```

```
Called_From(y.access_function, the_rm_ar_aux)

intra_rm_2: CONSTRAINT =
  (FORALL y: rm_main)
    [Called_From(the_rm_xa_aux.register, y)
     AND Called_From(the_rm_xa_aux.unregister, y)
     AND Called_From(y.close, the_rm_xa_aux)
     AND Called_From(y.commit, the_rm_xa_aux)
     AND Called_From(y.complete, the_rm_xa_aux)
     AND Called_From(y.end, the_rm_xa_aux)
     AND Called_From(y.forget, the_rm_xa_aux)
     AND Called_From(y.open, the_rm_xa_aux)
     AND Called_From(y.prepare, the_rm_xa_aux)
     AND Called_From(y.recover, the_rm_xa_aux)
     AND Called_From(y.rollback, the_rm_xa_aux)
     AND Called_From(y.start, the_rm_xa_aux)]]

END x_open_manager_alt_decomposition
```

```

%%% Just like 14, but an extra layer of abstraction in the RM to simplify
%%% the mapping.

x_open_manager_decomposition_2: ARCHITECTURE [ -> ]

IMPORTING ALL FROM X_Open_style

BEGIN

n: NAT      % Number of resource managers, a parameter in the specification
ar_requests, ar_resources: TYPE

COMPONENTS

ap: TYPE <= ARCHITECTURE [ -> ]

rm_ar_aux: TYPE <= { m: ARCHITECTURE [ -> ]
                      EXPORTING ALL
                      BEGIN
                        access_function: PROCEDURE [in: qt -> out: rt]
                      END m
                      | qt < ar_requests AND rt < ar_resources }

rm_xa_aux: TYPE <= ARCHITECTURE [ -> ]
                      EXPORTING ALL
                      BEGIN
                        close: XA_Close_Procedure
                            [info: XA_Info, rmid: INT, flags: INT
                             -> ret: INT]
                        commit: XA_Commit_Procedure
                            [id: X_Id, rmid: INT, flags: INT
                             -> ret: INT]
                        complete: XA_Complete_Procedure
                            [hdl: INT, retval: INT,
                             rmid: INT, flags: INT
                             -> ret: INT]
                        end: XA_End_Procedure
                            [id: X_Id, rmid: INT, flags: INT
                             -> ret: INT]
                        forget: XA_Forget_Procedure
                            [id: X_Id, rmid: INT, flags: INT
                             -> ret: INT]
                        open: XA_Open_Procedure
                            [info: XA_Info, rmid: INT, flags: INT
                             -> ret: INT]
                        prepare: XA_Prepares_Procedure
                            [id: X_Id, rmid: INT, flags: INT
                             -> ret: INT]
                        recover: XA_Recover_Procedure
                            [ids: X_Ids, count: INT,
                             rmid: INT, flags: INT
                             -> ret: INT]
                        rollback: XA_Rollback_Procedure
                            [id: X_Id, rmid: INT, flags: INT
                             -> ret: INT]
                        start: XA_Start_Procedure
                            [id: X_Id, rmid: INT, flags: INT
                             -> ret: INT]
                      END rm_xa_aux

rm_main: TYPE <= { m: ARCHITECTURE [ -> ]
                      EXPORTING ALL
                      BEGIN

```

```

access_function: RPC [in: qt -> out: rt]
close: RPC
    [info: XA_Info, rmid: INT, flags: INT
     -> ret: INT]
commit: RPC
    [id: X_Id, rmid: INT, flags: INT
     -> ret: INT]
complete: RPC
    [hdl: INT, retval: INT,
     rmid: INT, flags: INT
     -> ret: INT]
end: RPC
    [id: X_Id, rmid: INT, flags: INT
     -> ret: INT]
forget: RPC
    [id: X_Id, rmid: INT, flags: INT
     -> ret: INT]
open: RPC
    [info: XA_Info, rmid: INT, flags: INT
     -> ret: INT]
prepare: RPC
    [id: X_Id, rmid: INT, flags: INT
     -> ret: INT]
recover: RPC
    [ids: X_Ids, count: INT,
     rmid: INT, flags: INT
     -> ret: INT]
rollback: RPC
    [id: X_Id, rmid: INT, flags: INT
     -> ret: INT]
start: RPC
    [id: X_Id, rmid: INT, flags: INT
     -> ret: INT]
END m
| qt < ar_requests AND rt < ar_resources }

rm: TYPE <= ARCHITECTURE [ -> ]

rms: TYPE <= ARCHITECTURE [ -> ]

tm_aux: TYPE <= ARCHITECTURE [ -> ]
EXPORTING ALL
BEGIN
    begin: TX_Begin_Procedure [ -> ret: INT]
    close: TX_Close_Procedure [ -> ret: INT]
    commit: TX_Commit_Procedure [ -> ret: INT]
    information: TX_Info_Procedure [info: TX_Info -> ret: INT]
    open: TX_Open_Procedure [ -> ret: INT]
    rollback: TX_Rollback_Procedure [ -> ret: INT]
END tm_aux

tm_main: TYPE <= ARCHITECTURE [ -> ]
EXPORTING ALL
BEGIN
    register: AX_Register_Procedure
        [id: X_Id, rmid: INT, flags: INT
         -> ret: INT]
    unregister: AX_Unregister_Procedure
        [rmid: INT, flags: INT
         -> ret: INT]
    begin: RPC [ -> ret: INT]
    close: RPC [ -> ret: INT]
    commit: RPC [ -> ret: INT]
    information: RPC [info: TX_Info -> ret: INT]

```

```

        open: RPC [ -> ret: INT]
        rollback: RPC [ -> ret: INT]
    END tm_main

the_ap: ap
the_rms: rms
the_tm: tm
the_tm_main: tm_main
the_tm_aux: tm_aux

CONFIGURATION

rms_contents: CONSTRAINT =
    (FORALL y: COMPONENT)
        [y PROPERLY_CONTAINED_IN the_rms => (EXISTS z: rm) y CONTAINED_IN z]

rm_contents: CONSTRAINT =
    (FORALL y: COMPONENT)
        [(EXISTS z: rm) y PROPERLY_CONTAINED_IN z
         => (EXISTS z: rm_main) y CONTAINED_IN z
            OR (EXISTS z: rm_ar_aux) y CONTAINED_IN z
            OR (EXISTS z: rm_xa_aux) y CONTAINED_IN z]

rm_location: CONSTRAINT =
    (FORALL y: rm) y CONTAINED_IN the_rms

rm_main_location_1: CONSTRAINT =
    (FORALL y: rm_main) (EXISTS z: rm) y CONTAINED_IN z
rm_main_location_2: CONSTRAINT =
    (FORALL z: rm) (EXISTS y: rm_main) y CONTAINED_IN z

rm_ar_aux_location_1: CONSTRAINT =
    (FORALL y: rm_ar_aux) (EXISTS z: rm) y CONTAINED_IN z
rm_ar_aux_location_2: CONSTRAINT =
    (FORALL z: rm) (EXISTS y: rm_ar_aux) y CONTAINED_IN z

rm_xa_aux_location_1: CONSTRAINT =
    (FORALL y: rm_xa_aux) (EXISTS z: rm) y CONTAINED_IN z
rm_xa_aux_location_2: CONSTRAINT =
    (FORALL z: rm) (EXISTS y: rm_xa_aux) y CONTAINED_IN z

ar: CONSTRAINT =
    (FORALL u: rm_ar_aux) Called_From(u.access_function, the_ap)

tx: CONSTRAINT =
    Called_From(the_tm_aux.begin, the_ap)
    AND Called_From(the_tm_aux.close, the_ap)
    AND Called_From(the_tm_aux.commit, the_ap)
    AND Called_From(the_tm_aux.information, the_ap)
    AND Called_From(the_tm_aux.open, the_ap)
    AND Called_From(the_tm_aux.rollback, the_ap)

xa: CONSTRAINT =
    (FORALL v: rm_xa_aux)
        [Called_From(the_tm.register, v)
         AND Called_From(the_tm.unregister, v)
         AND Called_From(v.close, the_tm)
         AND Called_From(v.commit, the_tm)
         AND Called_From(v.complete, the_tm)
         AND Called_From(v.end, the_tm)
         AND Called_From(v.forget, the_tm)
         AND Called_From(v.open, the_tm)
         AND Called_From(v.prepare, the_tm)]

```

```

        AND Called_From(v.recover, the_tm)
        AND Called_From(v.rollback, the_tm)
        AND Called_From(v.start, the_tm) ]

intra_tm: CONSTRAINT =
    Called_From(the_tm_main.begin, the_tm_aux)
    AND Called_From(the_tm_main.close, the_tm_aux)
    AND Called_From(the_tm_main.commit, the_tm_aux)
    AND Called_From(the_tm_main.information, the_tm_aux)
    AND Called_From(the_tm_main.open, the_tm_aux)
    AND Called_From(the_tm_main.rollback, the_tm_aux)

intra_rm_1: CONSTRAINT =
    (FORALL y: rm_main) (EXISTS u: rm_ar_aux)
        Called_From(y.access_function, u)

intra_rm_2: CONSTRAINT =
    (FORALL u: rm_ar_aux) (EXISTS y: rm_main)
        Called_From(y.access_function, u)

intra_rm_3: CONSTRAINT =
    (FORALL y: rm_main) (EXISTS v: rm_xa_aux)
        [Called_From(v.register, y)
        AND Called_From(v.unregister, y)
        AND Called_From(y.close, v)
        AND Called_From(y.commit, v)
        AND Called_From(y.complete, v)
        AND Called_From(y.end, v)
        AND Called_From(y.forget, v)
        AND Called_From(y.open, v)
        AND Called_From(y.prepare, v)
        AND Called_From(y.recover, v)
        AND Called_From(y.rollback, v)
        AND Called_From(y.start, v) ]

intra_rm_4: CONSTRAINT =
    (FORALL v: rm_xa_aux) (EXISTS y: rm_main)
        [Called_From(v.register, y)
        AND Called_From(v.unregister, y)
        AND Called_From(y.close, v)
        AND Called_From(y.commit, v)
        AND Called_From(y.complete, v)
        AND Called_From(y.end, v)
        AND Called_From(y.forget, v)
        AND Called_From(y.open, v)
        AND Called_From(y.prepare, v)
        AND Called_From(y.recover, v)
        AND Called_From(y.rollback, v)
        AND Called_From(y.start, v) ]

END x_open_manager_decomposition_2

```

B SRI Publications: Correct Architecture Refinement

Correct Architecture Refinement

Mark Moriconi, Xiaolei Qian, and R. A. Riemenschneider

Abstract—A method is presented for the stepwise refinement of an abstract architecture into a relatively correct lower-level architecture that is intended to implement it. A refinement step involves the application of a predefined refinement pattern that provides a routine solution to a standard architectural design problem. A pattern contains an abstract architecture schema and a more detailed schema intended to implement it. The two schemas usually contain very different architectural concepts (from different architectural styles). Once a refinement pattern is proven correct, instances of it can be used without proof in developing specific architectures. Individual refinements are compositional, permitting incremental development and local reasoning. A special correctness criterion is defined for the domain of software architecture, as well as an accompanying proof technique. A useful syntactic form of correct composition is defined. The main points are illustrated by means of familiar architectures for a compiler. A prototype implementation of the method has been used successfully in a real application.

Keywords—Software architecture, hierarchy, stepwise refinement, refinement patterns, formal methods, relative correctness, composition

I. INTRODUCTION

DECISIONS about the architecture of a software system can have a major impact on system efficiency, maintainability, and evolvability. Architectural decisions typically are documented in terms of the ubiquitous box-and-arrow diagrams. Practicing engineers interpret the diagrams with respect to common architectural styles, such as dataflow, pipe-and-filter, batch-sequential, blackboard, implicit invocation (event-based), and client-server.

For a large system, its architecture often is described by a hierarchy of related architectures. An architecture hierarchy is a linear sequence of two or more individual architectures that may differ with respect to the number and kind of components and connections among them. For example, an abstract architecture containing functional components related by dataflow connections may be implemented in a concrete architecture in terms of procedures, control connections, and shared variables. In general, an abstract architecture is smaller and easier to understand; a concrete architecture reflects more implementation concerns.

The utility of an architecture hierarchy is severely limited by the current level of informality. Individual architectures may be ambiguous, allowing multiple and perhaps unintended interpretations. The mapping between architectures in the hierarchy is partially specified, if at all, making it impossible to accurately trace the lineage of implementation decisions. The analysis of architecture is limited to

This research was supported in part by the Advanced Research Projects Agency under Rome Laboratory contract F30602-93-C-0245.

The authors are with the Computer Science Laboratory, SRI International, Menlo Park, California 94025. Email: {moriconi, qian, rar}@csli.sri.com.

syntactic checks. It is not possible to check semantic properties of an architecture, such as the safety and fairness of its connections, or to check the relative correctness of two architectures in the hierarchy. Consequently, a concrete architecture may erroneously be seen as an implementation of a more abstract architecture.

The main contribution in this paper is a methodology for the correct stepwise refinement of software architectures. It is expected to lead to fewer architectural design errors, to extensive and systematic reuse of design knowledge and proofs, and ultimately to an architecture synthesis tool similar to those now used for integrated circuit design. The methodology involves the use of instances of architecture refinement patterns that are correctness preserving and compositional.

A refinement pattern provides a routine solution to a standard architectural design problem. For example, a pattern may show how to implement a single dataflow connection in shared memory, or several patterns may combine to implement dataflow diagrams in terms of some form of client/server architecture. A pattern contains a pair of architecture schemas that are proved to be relatively correct with respect to a given mapping schema between them. The proof is performed only once; every instance of a refinement pattern is guaranteed to be correct. A schema can be homogeneous (consisting of one style) or heterogeneous (consisting of multiple styles). The two schemas in a refinement pattern may, and usually do, contain concepts from different architectural styles.

A useful form of correctness-preserving composition is defined that applies to both individual refinements and existing architectures. The latter is important because we want to be able to assemble existing subsystem architectures into a single system. Two architectures can be composed even if their vocabularies are not disjoint. In general, "horizontal" composition requires a case-by-case proof of correctness. However, we define a simple syntactic criterion that, if satisfied, guarantees compositionality. Because our correctness relation is transitive, the "vertical" composition of levels in an architecture hierarchy preserves correctness, and we are guaranteed that the most concrete architecture in the hierarchy meets the requirements of the most abstract architecture in the hierarchy.

The correctness of architecture refinement and composition involves a special correctness criterion, which is stronger than the usual one for functional refinement, and a special mapping between architectures, that is more complex than the usual mapping between data structures. A mapping between architectures involves an extensive translation in which the representation of components, interfaces, and connections may change and, moreover, these abstract objects may be aggregated, decomposed, or elim-

inated in the concrete architecture.

A stronger correctness criterion is needed because of the potential uses of architectures. Consider the role an architecture can play in reducing the time to provide fixes, optimizations, and upgrades to systems in deployment. If the architecture accurately models the implementation, it can be used to focus and explore the consequences of changes to the implementation. But if the implementation contains connections that do not appear in the architecture, a developer could easily be misled into making changes that appear to be minor and localized but that, in fact, have widespread consequences. For example, we may specify a pipeline architecture, restricting the system topology to a linear sequence of filters, to facilitate component reusability. If the concrete architecture implements the pipeline, but additionally introduces feedback loops, the *raison d'être* behind the original pipeline architecture is no longer valid. In general, the preservation of "communication integrity" is integral to the utility of an architecture.

Therefore, an architecture should describe explicitly the components, interfaces, and connections that are required of the target system, and perhaps more importantly, those that are *not* intended to appear in the target system. This observation leads to a *completeness assumption* about a given architecture, namely that an architecture contains *all* components, interfaces, and connections intended to be true of the architecture at its level of detail. If a fact is not explicit in the architecture, or deducible from it, we assume that it is not intended to be true of the architecture. In the pipeline example, we couple the linearity property with the completeness assumption to infer that no feedback loop is allowed in an implementation of the architecture. In general, an architecture (whether static or dynamic) can contain an unbounded number of facts.

The completeness assumption requires that we prove not only that a concrete architecture does not lose properties of the abstract architecture, but also that no new properties about the abstract architecture can be inferred from the concrete architecture. The standard method for reasoning about the relative correctness of two specifications is to show that the concrete specification logically implies the abstract specification under a given mapping between them. This allows an implementation to exhibit additional, unspecified behaviors, as long as the specified behavior is implemented. If the standard proof method is applied to architectures, there would be no guarantee that negative properties are preserved under refinement.

Fortunately, there is a well-understood mathematical property, called *faithful interpretation*, that can be adapted for our purposes. If a certain mapping between the two architectures is faithful, both the positive and the implicit negative facts in the abstract architecture are preserved in the concrete architecture. However, a proof of faithfulness is inherently hard, and we are not aware of any general proof technique in the literature. We introduce a systematic technique for proving faithfulness. The inherent complexity of such proofs is one reason why we advocate a methodology that makes use of preproved refinement pat-

terns.

It is worth mentioning that an important consequence of the completeness assumption is that the standard stepwise refinement paradigm is unsound with respect to the correctness relation. Certain refinements of an architecture must be composed horizontally. Completed levels in an architecture hierarchy can be composed vertically.

This paper is organized as follows. The next section illustrates the refinement problem and our approach to a solution. Section III makes useful distinctions among architectural styles, architecture schemas, and instance architectures, and shows how they can be represented as logical theories. We use first-order theories, but our basic framework does not depend on a particular logic. By formalizing architectures and their properties in logic, our results can be applied to a large class of architecture definition languages. Sections IV, V, and VI discuss mappings, correctness, and composition, respectively.

Section VII presents several different refinement patterns that are used in Section VIII in the development of standard architectures for a compiler. The development includes both refinement and composition. Sections IX reports on a larger experiment involving an operational power-control system. Section X describes related work, and the last section summarizes our results, their implications, and makes suggestions for future work.

II. ILLUSTRATION OF APPROACH TO REFINEMENT

A software architecture is represented using the following concepts.

1. **Component:** An object with independent existence, e.g., a module, process, procedure, or variable.
2. **Interface:** A typed object that denotes a logical point of interaction between a component and its environment.
3. **Connector:** A typed object relating interface points, components, or both.
4. **Configuration:** A collection of constraints that wire objects into a specific architecture.
5. **Mapping:** A relation that defines a syntactical translation from the language of an abstract architecture to the language of a concrete architecture.
6. **Architectural style:** For the purposes of this paper, a style consists of a vocabulary of design elements, well-formedness constraints that determine how they can be used, and a semantic definition of the connectors associated with the style.

Components, interfaces, and connectors are treated as *first-class objects* — i.e., they have a name and they are refinable. Abstract architectural objects can be decomposed, aggregated, or eliminated in a concrete architecture. The semantics of components is not considered part of an architecture, but the semantics of connectors is.

Consider the standard dataflow architecture for a compiler that is depicted at the top of Figure 1. The diagram is intended to convey an intuitive feel for the architecture; it is not a formal description of the architecture. Boxes denote functional components and arrows denote direc-




Fig. 1. Two architectures for a compiler

tional dataflow between ports. The labels on arrows denote types or value domains. A value cannot be transmitted between ports unless its type is compatible with the types of the ports. By the completeness assumption, this dataflow model of the compiler fixes its functional units, their interfaces, and the direction, source, and destination of all of its flows.

A textual specification of the dataflow architecture is contained in Figure 2. A dataflow component is a function with a signature describing its interface. Four dataflow connectors are declared to carry values of various types. The configuration assertions wire the connectors and interfaces together into a specific type-consistent architecture. The module imports various types and the functional and dataflow styles for use in the specification of the architecture.

A concrete architecture intended to implement the dataflow model of the compiler is depicted at the bottom of Figure 1. The concrete architecture is a hybrid that implements the dataflow style in terms of pipe-filter, batch-sequential, and shared-memory styles. Abstract signatures have been changed, dataflow connectors have been implemented in several ways, new components (data objects) are introduced, and precedence relations are added to preserve the original flows in the presence of shared-memory communication.¹ A textual specification of the level-2 architecture of the compiler can be found in the appendix.

We do not want to construct the level-1 and the level-2 architectures and then perform an after-the-fact correctness proof. Instead, we want to systematically and incrementally transform the level-1 architecture into the level-2 architecture. The level-2 architecture should be correct by construction, requiring no explicit proofs in its derivation. This can be accomplished through a series of small, local refinements, each of which involves the application of a correct refinement pattern. Then, the local refinements are combined to form the larger composite level-2 architecture, which is guaranteed to correctly implement the level-1

¹A dataflow connection is treated as an intransitive relation.

```

compiler_L1: MODULE
  [char_iport: SEQ(character) -> code_oport: code]
  IMPORT character, code, token, binding, ast
  FROM compiler_types
  IMPORT Function FROM Functional_Style
  IMPORT Dataflow_Channel, Connects
  FROM Dataflow_Style

COMPONENTS
  lexical_analyzer: Function
    [char_iport: SEQ(character)
     -> token_oport: SEQ(token),
      bind_oport: SEQ(binding)]
  parser: Function
    [token_iport: SEQ(token)
     -> base_ast_oport: ast]
  analyzer_optimizer: Function
    [base_ast_iport: ast, bind_iport: SEQ(binding)
     -> full_ast_oport: ast]
  code_generator: Function
    [full_ast_iport: ast -> code_oport: code]

CONNECTORS
  token_channel: Dataflow_Channel[SEQ(token)]
  bind_channel: Dataflow_Channel[SEQ(binding)]
  base_ast_channel: Dataflow_Channel[ast]
  full_ast_channel: Dataflow_Channel[ast]

CONFIGURATION
  token_flow:
    Connects(token_channel, token_oport, token_iport)
  bind_flow:
    Connects(bind_channel, bind_oport, bind_iport)
  base_ast_flow:
    Connects(base_ast_channel,
             base_ast_oport, base_ast_iport)
  full_ast_flow:
    Connects(full_ast_channel,
             full_ast_oport, full_ast_iport)

END compiler_L1

```

Fig. 2. Specification of dataflow architecture for the compiler

architecture.

As an illustration of our approach, consider the implementation of the dataflow channel between the parser and analyzer in terms of the reading and writing of a shared abstract syntax tree. More specifically, we propose to refine abstract subarchitecture

```

parser:           Function [-> base_ast_oport: ast]
analyzer_optimizer: Function [base_ast_iport: ast -> ]
base_ast_channel: Dataflow_Channel[ast]
base_ast_flow:
  Connects(base_ast_channel,
           base_ast_oport, base_ast_iport)

```

into concrete subarchitecture

```

parser:           Function[->]
analyzer_optimizer: Function[->]
abstract_syntax_tree: Variable[ast]
write_base_ast: Writes(parser, abstract_syntax_tree)
read_base_ast:
  Reads(analyzer_optimizer, abstract_syntax_tree)

```

For simplicity, the component signatures contain only the ports that are relevant to this refinement. The dataflow connection is implemented by a component (a shared variable containing the tree) and two connections (the read

PATTERN OF ABSTRACT ARCHITECTURE:
<i>M</i> : MODULE[→]
COMPONENTS
<i>f</i> ₁ : Functional_Style!Function[→ <i>op</i> : <i>t</i>]
<i>f</i> ₂ : Functional_Style!Function[<i>ip</i> : <i>t</i> →]
CONNECTORS
<i>c</i> : Dataflow_Style!Dataflow_Channel[<i>t</i>]
CONFIGURATION
<i>a</i> : Dataflow_Style!Connects(<i>c</i> , <i>op</i> , <i>ip</i>)
PATTERN OF CONCRETE ARCHITECTURE:
<i>M</i> : MODULE[→]
COMPONENTS
<i>f</i> ₁ : Functional_Style!Function[→]
<i>f</i> ₂ : Functional_Style!Function[→]
<i>m</i> : Shared_Memory_Style!Variable[<i>t</i>]
CONFIGURATION
<i>a</i> ₁ : Shared_Memory_Style!Writes(<i>f</i> ₁ , <i>m</i>)
<i>a</i> ₂ : Shared_Memory_Style!Reads(<i>f</i> ₂ , <i>m</i>)
ABSTRACT TO CONCRETE ASSOCIATIONS:
<i>op</i> →→ <i>c</i> →→ <i>m</i> <i>a</i> →→ (<i>a</i> ₁ , <i>a</i> ₂)
<i>ip</i> →→

Fig. 3. Simple refinement pattern

and write relations).² The new concrete signature for the parser and the analyzer reflects the difference between port-to-port communication and direct shared-memory communication through a variable. As an analogous example, consider an architecture consisting of two procedures that communicate solely by means of procedure calls. If we optimize this architecture so that large objects are no longer transmitted by value, but instead are accessed directly as shared objects, the signatures of the two procedures would change.

The refinement pattern in Figure 3 specifies a way to implement dataflow in terms of the reading and writing of a single variable. The read and write relations in the concrete schema are primitives that cannot be refined. The italic letters denote schema variables that can be instantiated with object names, and the symbol “!” is used to qualify names. The pattern can be proven correct with respect to the four associations shown at the bottom of the pattern.³

The abstract schema in the pattern matches the level-1 subarchitecture. However, if the same substitutions are made in the concrete schema, three schema variables are left uninstantiated — namely, *m*, *a*₁, and *a*₂. Of course, any unused names could be substituted. Let us assume that the architect selects mnemonic names that give the following associations.

```
base_ast_oport   →→
base_ast_iport   →→
base_ast_channel →→ abstract_syntax_tree
base_ast_flow    →→ (write_base_ast, read_base_ast)
```

²The shared abstract syntax tree could have been represented as an encapsulated data type. If we had chosen that representation, the architecture would involve calls to access functions that read and write the internal variable used to represent the tree.

³In a correctness proof, the associations in the pattern are incorporated into a more complex mapping between the first-order theories that represent the abstract and concrete architectures.

Since this instance of the pattern matches the abstract subarchitecture of the compiler and since all instances of the pattern are guaranteed to preserve correctness, we can conclude that the proposed refinement is correct.

In a later section, we define enough patterns to transform the full level-1 compiler architecture into the full level-2 architecture. Additional patterns are defined that can be used to transform the level-2 architecture into a more efficient batch-sequential architecture. The final batch-sequential architecture can be found in the appendix. The completed compiler architecture can be connected to other subsystem architectures, such as the file system architecture, to form a correct composite system.

III. ARCHITECTURES AS THEORIES

We want to leave open the choice of language for specifying an architecture. Therefore, we will represent architectures as logical theories. We find it convenient to use first-order theories; however, our results do not depend on this choice.

It is useful to distinguish among three related architectural theories:

- An *architectural style* is a theory consisting of a vocabulary of the relevant architectural concepts and well-formedness axioms that determine how they can be used. Also associated with a style are rules for translating textual specifications in the style into their underlying logical representation.
- An *architecture* is a theory consisting of one or more style subtheories and possibly an infinite number of constants that are names of the objects in the particular architecture. The axioms of the theory are the style axioms and possibly additional axioms that relate the constants.
- An *architecture schema* is an architecture containing one or more schema variables. An *instance* of an architecture schema is obtained by substitution of constants for all of its schema variables. An instance of an architecture schema is sometimes called an *instance architecture* or an *instance theory*.

A. Architectural Styles

Consider the dataflow style. Its vocabulary contains predicates for describing functional components, ports, values associated with ports, dataflow channels, values associated with dataflow channels, and connections of channels to ports. More precisely, the following sorts denote the first-class objects in a dataflow theory: *channel*, *function*, *iport*, and *oport*. We also make use of sorts *bool* and *val*, where *val* denotes the set of all possible values. The dataflow style has the following operations.

```
OutPort: oport × function → bool
Supplies: oport × val → bool
InPort: iport × function → bool
Accepts: iport × val → bool
Carries: channel × val → bool
Connects: channel × oport × iport → bool
```

These predicates are used to represent a dataflow architecture in ordinary first-order logic. Sorts can be represented as unary predicates but, for simplicity, we omit them in formulas.

An example of a well-formedness axiom is that every function must have at least one port:

$$\forall x \exists y [\text{InPort}(y, x) \vee \text{OutPort}(y, x)]$$

Another requirement is that a channel attached to an output port must be able to carry any value supplied by the port:

$$\begin{aligned} & \forall x \forall y [\exists z \text{Connects}(x, y, z) \\ & \quad \supset \forall v [\text{Supplies}(y, v) \supset \text{Carries}(x, v)]] \end{aligned}$$

B. Translation to Logic

Architectures and refinement patterns are expressed in a readable textual language. To reason about them, they are translated into logic by means of simple “theory generation rules” which are associated with architectural styles. For the dataflow style, if the specification of an architecture contains an instance of function declaration schema

$$f: \text{Functional_Style!Function}[\rightarrow op: t]$$

the underlying theory contains the same instance of first-order sentences

$$\begin{aligned} & \text{OutPort}(op, f) \\ & \forall v [\text{Supplies}(op, v) \supset t(v)] \end{aligned}$$

Similarly, a function declaration of the form

$$f: \text{Functional_Style!Function}[ip: t \rightarrow]$$

is translated to axioms

$$\begin{aligned} & \text{InPort}(ip, f) \\ & \forall v [t(v) \supset \text{Accepts}(ip, v)] \end{aligned}$$

Dataflow connector

$$c: \text{Dataflow_Style!Dataflow.Channel}[t]$$

translates to

$$\forall v [t(v) \supset \text{Carries}(c, v)]$$

and configuration constraint

$$a: \text{Dataflow_Style!Connects}(c, op, ip)$$

to

$$\text{Connects}(c, op, ip)$$

which is not an object and, therefore, is not named in the logic.

C. Architecture Schemas

The two schemas appearing in the pattern of Figure 3 will be referred to throughout the paper. Theory Θ_D corresponds to the abstract schema and theory Θ_M corresponds to the concrete schema.

Theory Θ_D is formed by applying the theory generation rules of the dataflow style to the abstract schema, which gives

$$\begin{aligned} & \text{OutPort}(op, f_1) \\ & \forall v [\text{Supplies}(op, v) \supset t(v)] \\ & \text{In.Port}(ip, f_2) \\ & \forall v [t(v) \supset \text{Accepts}(ip, v)] \\ & \forall v [t(v) \supset \text{Carries}(c, v)] \\ & \text{Connects}(c, op, ip) \end{aligned}$$

This theory satisfies the two well-formedness axioms stated earlier.

The concrete architecture schema in Figure 3 is written in a shared-memory style, which permits the reading and writing of a shared variable. Shared-variable communication is modeled using a call site as the interface between a function and the shared variable.⁴ A call site serves the same purpose as a port in the dataflow style. The name of every different call site must be unique. Θ_M has the following style-specific sorts: *variable* denotes the set of all possible variables and *site* denotes the set of all possible call sites of which there are two kinds. The sort *rsite* denotes the sites that read, or input, values; the sort *wsite* denotes the ones that write, or output, values. The signature for Θ_M is

$$\begin{aligned} & \text{Holds: } \text{variable} \times \text{val} \rightarrow \text{bool} \\ & \text{CallSiteOf: } \text{site} \times \text{function} \rightarrow \text{bool} \\ & \text{Writes: } \text{wsite} \times \text{variable} \rightarrow \text{bool} \\ & \text{Puts: } \text{wsite} \times \text{val} \rightarrow \text{bool} \\ & \text{Reads: } \text{rsite} \times \text{variable} \rightarrow \text{bool} \\ & \text{Gets: } \text{rsite} \times \text{val} \rightarrow \text{bool} \end{aligned}$$

The axioms of Θ_M are

$$\begin{aligned} & \forall v [t(v) \supset \text{Holds}(m, v)] \\ & \text{CallSiteOf}(w, f_1) \\ & \text{Writes}(w, m) \\ & \forall v [\text{Puts}(w, v) \supset t(v)] \\ & \text{CallSiteOf}(r, f_2) \\ & \text{Reads}(r, m) \\ & \forall v [t(v) \supset \text{Gets}(r, v)] \end{aligned}$$

which must satisfy the well-formedness axioms for the shared-memory style. Schema variables *r* and *w* denote names of call sites and do not appear in Figure 3.

IV. MAPPINGS

To prove the relative correctness of two architectures, we must specify a mapping between them. An *interpretation mapping* is an association between formulas of the language of the abstract theory and formulas of the language of the concrete theory. An interpretation mapping is determined using two different mappings.

- A *name mapping* associates the objects declared in an abstract architecture with objects declared in a concrete architecture.
- A *style mapping* says how the constructs of an abstract-level style can be implemented in terms of the constructs of a concrete-level style. More specifically,

⁴We could have chosen not to model call sites or some equivalent interface object, but this would require a more liberal definition of interpretation than the one given in this paper. The present model simplifies the mapping from Θ_D to Θ_M .

it maps uninstantiated predicates of the abstract-level language to uninstantiated formulas of the concrete-level language.

Style mappings can be complicated, but need to be defined and proved only once. Name mappings are much simpler and are specific to a given pair of architectures.

A name mapping is determined by the identifier associations in a given refinement pattern. For example, association $c \rightarrow m$ in Figure 3 says that channel c of the abstract schema is mapped to variable m of the concrete schema. Association $op \rightarrow$ says that the concrete object that corresponds to abstract port op is not explicitly named in the concrete schema. Since we have chosen a shared-memory model that has call sites corresponding to ports, we are free to introduce any unused name for the sites.

Let N_M^D be name mapping

$$\begin{array}{rcl} c & \mapsto & m \\ op & \mapsto & w \\ ip & \mapsto & r \end{array}$$

which relates objects in Θ_D to their refinements in Θ_M . Observe that not every association in the refinement pattern appears in the name mapping. Identifiers a , a_1 , and a_2 refer to part of the specification but do not name objects. Hence, they do not appear in the logical representation. The domain of a name mapping can be extended to include all abstract-level terms by mapping variables to themselves.⁵

Let S_M^D denote the general mapping from the dataflow style to the shared-memory style:

$$\begin{aligned} \text{Function}(_) &\mapsto \text{Function}(_) \\ \text{OutPort}(_, _) &\mapsto \text{CallSiteOf}(_, _) \wedge \exists v \text{ Puts}(_, v) \\ \text{Supplies}(_, _) &\mapsto \text{Puts}(_, _) \\ \text{InPort}(_, _) &\mapsto \text{CallSiteOf}(_, _) \wedge \exists v \text{ Gets}(_, v) \\ \text{Accepts}(_, _) &\mapsto \text{Gets}(_, _) \\ \text{Channel}(_) &\mapsto \text{Variable}(_) \\ \text{Carries}(_, _) &\mapsto \text{Holds}(_, _) \\ \text{Connects}(_, _, _) &\mapsto \text{Writes}(_, _) \wedge \text{Reads}(_, _) \end{aligned}$$

The Puts and Gets predicates ensure that the right kind of site is associated with each port.

The last association specifies the implementation strategy. In Θ_D we have $\text{Connects}(c, op, ip)$, which can be implemented by having the call that corresponds to op perform a write operation on the variable that corresponds to channel c , and the one that corresponds to ip read the variable that corresponds to c . The other associations say that channels are mapped to variables, that output ports are mapped to calls that supply values, and that input ports are mapped to calls that receive values.

An *interpretation mapping* I is determined from a name mapping N and a style mapping S , as follows: for every

⁵Note that our languages contain no function symbols. A formal treatment of interpretations for languages that include them can be found in [6].

predicate P , all terms t_1, t_2, \dots, t_n , every variable x , and all formulas F and G of the abstract language,

$$\begin{aligned} I(P(t_1, t_2, \dots, t_n)) &= S(P)(N(t_1), N(t_2), \dots, N(t_n)) \\ I(\neg F) &= \neg(I(F)) \\ I(F \wedge G) &= I(F) \wedge I(G) \\ I(F \vee G) &= I(F) \vee I(G) \\ I(F \supset G) &= I(F) \supset I(G) \\ I(\forall x F) &= \forall x I(F)^6 \\ I(\exists x F) &= \exists x I(F) \end{aligned}$$

Let I_M^D denote the interpretation mapping from theory Θ_D to theory Θ_M . Both the basic facts and the general well-formedness axioms in Θ_D must be mapped. For example,

$$\begin{aligned} I_M^D(\text{Connects}(c, op, ip)) &= S_M^D(\text{Connects})(N_M^D(c), N_M^D(op), N_M^D(ip)) \\ &= S_M^D(\text{Connects})(m, w, r) \\ &= \text{Writes}(w, m) \wedge \text{Reads}(r, m) \end{aligned}$$

which is the intended implementation. Similarly, the general dataflow-style requirement that each function have at least one input or output port maps to the shared-memory requirement that each function have a call site that can input or output values. That is,

$$\begin{aligned} I_M^D(\forall x \exists y [\text{InPort}(y, x) \vee \text{OutPort}(y, x)]) &= \forall x \exists y [I_M^D(\text{InPort}(y, x)) \vee I_M^D(\text{OutPort}(y, x))] \\ &= \forall x \exists y [S_M^D(\text{InPort})(N_M^D(y), N_M^D(x)) \\ &\quad \vee S_M^D(\text{OutPort})(N_M^D(y), N_M^D(x))] \\ &= \forall x \exists y [(\text{CallSiteOf}(y, x) \wedge \exists v \text{ Gets}(y, v)) \\ &\quad \vee (\text{CallSiteOf}(y, x) \wedge \exists v \text{ Puts}(y, v))] \end{aligned}$$

V. CORRECTNESS

Two instance architectures, represented as theories, are proven correct with respect to an interpretation mapping between them and the completeness assumption. An interpretation mapping contains a style mapping whose semantic correctness should be established as a proof obligation. Proof of style mappings is discussed in a companion paper [18], which gives a proof of mapping S_M^D from the dataflow to the shared-memory style. The connectors in the styles are defined in a temporal logic, and both safety and fairness conditions are shown to be satisfied by the shared-memory implementation. The safety condition is that the shared-memory implementation preserves order and does not lose values; the fairness condition is that all values written into shared memory will eventually be read. The proof of a style mapping is performed only once; it need not be repeated when the two styles are used.

A. Criterion

Let Θ and Θ' be instance theories (containing no schema variables) associated with an abstract and a concrete architecture, respectively. Let I be an interpretation mapping

⁶In general, the range of quantifiers must be restricted to a subset of the concrete domain, see [6]. But no restriction is required for our example, because every concrete-level object implements an abstract-level object.

from the language of Θ to the language of Θ' . For every sentence F , mapping I is a *theory interpretation* provided

$$\text{if } F \in \Theta \text{ then } I(F) \in \Theta'$$

This is the usual definition of correctness.

Since a given architecture is assumed to be complete with respect to its level of detail, we additionally require that the concrete architecture add no new facts about the abstract architecture. To prove this, we must additionally show that

$$\text{if } F \notin \Theta \text{ then } I(F) \notin \Theta'$$

in which case I is a *faithful interpretation*. This says that, if a sentence is not in the abstract theory, its image cannot be in the concrete theory. Observe that Θ' is a conservative extension of Θ provided the identity map faithfully interprets Θ in Θ' .

B. Proof Technique

Again, let Θ and Θ' be instance theories and I be the interpretation mapping between them. We present a general model-theoretic proof technique for showing that interpretation mapping I is a faithful interpretation of abstract theory Θ in concrete theory Θ' . First, we prove that I is a theory interpretation of Θ in Θ' . This can be done by means of a standard proof technique: *For every axiom in Θ , establish that the image of the axiom under I is a logical consequence of the axioms of Θ'* .

Second, we must prove that interpretation mapping I is a faithful. The proof method has to take into account that there is no direct method for determining that a formula is not in Θ' . Our proof technique for faithfulness is based on two model-theoretic concepts:

- The interpretation mapping I from Θ to Θ' induces a mapping I' from structures of the concrete language to structures of the abstract language.⁷ Given a structure \mathcal{A}' of the concrete language, I' maps \mathcal{A}' to a structure \mathcal{A} of the abstract language as follows. The universe of \mathcal{A} is the same as the universe of \mathcal{A}' . If I maps atomic formula $P(x_1, x_2, \dots, x_n)$ to concrete formula F , then I' assigns to predicate P in the abstract language the set of tuples in \mathcal{A}' that satisfy F .
- The theory that describes structure \mathcal{A} is obtained as follows. First, expand the language of \mathcal{A} to include a name for every member of the universe of \mathcal{A} . Next, expand \mathcal{A} by assigning every new name to the appropriate member of \mathcal{A} . The theory that describes \mathcal{A} is the set of sentences in the expanded language that are true in the expanded structure.

Our technique for proving the faithfulness of I can now be stated as follows: *For every model \mathcal{A} of Θ , find a model \mathcal{A}' of Θ' such that the image of \mathcal{A}' under the induced mapping I' can be expanded to a model of the theory that describes \mathcal{A}* . This model-theoretic characterization of faithfulness is equivalent to our theory-based definition of correctness.

⁷Recall from logic that a structure of a first-order language consists of a universe and the assignment of elements of the universe to the constants and relations over the universe to the function and predicate symbols.

Roughly speaking, this characterization requires that, for every model \mathcal{A} of Θ , there is a model \mathcal{A}' of Θ' such that \mathcal{A} and $I'(\mathcal{A}')$ cannot be distinguished using the resources of first-order logic. If we were to use an architectural specification language based on some other logic, a similar characterization based on the expressive power of that logic would be substituted. For example, if the content of our architectural specifications were expressed in type theory, we would require that $I'(\mathcal{A}')$ can be expanded to model every type-theoretic sentence expressible in the language that contains a name for every object in the domain of \mathcal{A} , every relation among those objects, every relation among those relations, and so on, that is true in \mathcal{A} . (It is easy to see that this amounts to requiring that $I'(\mathcal{A}')$ and \mathcal{A} be isomorphic.) So our general method for demonstrating faithfulness can be used with any logic-based architectural specification language, as long as the question of whether a structure that represents an architecture satisfies a specification has a well-defined answer.

C. Application to Refinement Patterns

A refinement pattern consists of a triple (Θ, Θ', N) where Θ and Θ' are theories containing schema variables and N is a name mapping from Θ to Θ' . A pattern is correct provided every instance of Θ and Θ' is relatively correct with respect to the same instance of interpretation mapping $I : \Theta \rightarrow \Theta'$ determined by mapping N and the relevant style mapping(s).

Consider theories Θ_D and Θ_M related by interpretation mapping I_M^D . We must show that, for every instantiation of the schema variables, I_M^D is a theory interpretation of Θ_D in Θ_M and I_M^D is faithful. The former is straightforward.

To prove faithfulness, consider the induced mapping of I_M^D . If M is a structure for Θ' , then the induced mapping applied to M is a structure D for the dataflow language. The only interesting assignment is to the predicate *Connects*, which is the set of tuples

$$\{(x, y, z) \in |M|^3 : M \models \text{Writes}(y, x) \wedge \text{Reads}(z, x)\}$$

because I_M^D maps *Connects*(c, op, ip) to the formula

$$\text{Writes}(w, m) \wedge \text{Reads}(r, m)$$

where c, op, ip, w, m , and r are schema variables.

To show that I_M^D is faithful, we use I_M^D to transform a model D of an instance of Θ_D to a model M of an instance of Θ_M . The universe of M is the same as D in this example. The predicate *Function* is assigned to the set of all objects that are functions in D , namely,

$$\{x \in |D| : D \models \text{Function}(x)\}$$

so that D and M agree on functions. The predicate *Variable* is assigned to

$$\{x \in |D| : D \models \text{Channel}(x)\},$$

the predicate *Reads* is assigned to

$$\{(x, y) \in |D|^2 : \text{for some } z \text{ in } |D|, D \models \text{Connects}(y, z, x)\}$$

and similarly for the remaining predicates. The image of \mathcal{M} under the induced mapping is \mathcal{D} . Obviously, \mathcal{D} can be expanded to a model of the theory that describes \mathcal{D} . Therefore, $I_{\mathcal{M}}^P$ is faithful. Note that, since the image of \mathcal{M} under the induced mapping is identical to \mathcal{D} , the interpretation $I_{\mathcal{M}}^P$ would remain faithful if we were to switch from first-order logic to some stronger logic, such as type theory.

VI. COMPOSITION

We define two forms of composition for instance architectures. Horizontal composition is used to compose instances of refinement patterns to form one large composite refinement architecture. It is also used to compose existing architectures into larger architectures. Vertical composition is used to chain together a sequence of correct architectures, allowing us to conclude that the most concrete architecture in a hierarchy is correct with respect to the most abstract architecture in the hierarchy. Vertical composition is justified since faithful interpretation is transitive.

Let Θ_1 and Θ_2 be instance theories that represent two abstract architectures. Let Θ'_1 and Θ'_2 be concrete theories intended to implement Θ_1 and Θ_2 , respectively. Two pairs of architecture theories can be composed only in ways that preserve faithfulness. More precisely, if

$$I_1: \Theta_1 \rightarrow \Theta'_1 \text{ and } I_2: \Theta_2 \rightarrow \Theta'_2$$

are faithful interpretations, then we want

$$I_1 \cup I_2: \Theta_1 \cup \Theta_2 \rightarrow \Theta'_1 \cup \Theta'_2$$

to be a faithful interpretation. (The union of two theories is the deductive closure of the set-theoretic union of the theories.)

This property holds provided two general conditions are satisfied.

1. The composite interpretation mapping must be a function. For a sentence F , we require that

$$\text{if } F \in \Theta_1 \cap \Theta_2 \text{ then } I_1(F) = I_2(F)$$

which guarantees that interpretation mappings I_1 and I_2 agree on shared objects and shared style constructs.

2. It must not be possible to infer new facts about the composite abstract architecture from the composite concrete architecture. That is, for language L_1 of Θ_1 and L_2 of Θ_2 , if

$$F \text{ is a sentence of } L_1 \cup L_2$$

and

$$\Theta'_1 \cup \Theta'_2 \vdash I_1 \cup I_2(F)$$

then we must prove that

$$I_1[\Theta_1] \cup I_2[\Theta_2] \vdash I_1 \cup I_2(F).$$

The intuition behind the second condition can be illustrated by means of a simple example. Consider an architecture in which there is a dataflow connection from

A to B and another architecture that has dataflow connection from B to C . Suppose that both flows are implemented correctly in concrete architectures, but that in one A writes some variable x and in the other C reads a variable x . Each implementation is correct, since neither introduces a new dataflow. However, the composite concrete architecture reads and writes x , from which we can infer an entirely new abstract dataflow connection from A to C . Consequently, the composite abstract architecture is not faithfully interpreted (by the composite mapping) in the composite concrete architecture (under the original assumption that dataflow is intransitive).

Of course, we do not want to have to prove that every refinement pattern can be composed with every other refinement pattern. Instead, we would like simple syntactic criterion that, if satisfied, guarantees compositionality. One such criterion is that the two abstract architectures can share only components and lower-level architectures can share only images of those components under the interpretation mapping. This means that an architecture cannot contain certain global assertions, such as a requirement that there are exactly three connections in any architecture.

An example of the horizontal composition of pattern instances involves the compiler architecture in Figure 1. We have proved that the dataflow connection between the parser and the analyzer is implemented correctly by means of the reading and writing of the tree, using instances of Θ_D , Θ_M , and $I_{\mathcal{M}}^P$ from Figure 3. Similarly, we can show that the dataflow connection from the lexical analyzer to the parser is correctly implemented by the pipeline connection. The two architectures share only one component, the parser. Therefore, our second condition is satisfied and we can compose them without further proof.

A different kind of example is contained in Figure 4. We want to compose two architectures, called "subsystem A" and "subsystem B", into a single system architecture. We construct a new architecture with components "A" and "B" connected through new interfaces. According to our syntactic constraint, the three architectures can be combined to form a composite system that is correct if the three subsystems are.

VII. SOME REFINEMENT PATTERNS

We present six broadly useful patterns for refining components, connectors, and interfaces.⁸ The patterns involve several common architecture styles and each pattern has been proven correct.

A refinement pattern is presented in a table containing two architecture schemas, an association of abstract and concrete objects, and possibly constraints on one or both of the schemas. By convention, a schema variable that occurs in both an abstract and a concrete schema must match the same object, modulo renaming. We prime a concrete schema variable to indicate that it is the name of a new object not associated with any abstract-level object, or that it

⁸Type refinement is not covered because it requires a somewhat different correctness criterion.




Fig. 4. Illustration of Subsystem Composition

Symbol	Style Name
BS	Batch_Sequential_Style
CT	Control_Transfer_Style
D	Dataflow_Style
F	Functional_Style
PP	Process_Pipeline_Style
SM	Shared_Memory_Style

TABLE I
ABBREVIATIONS FOR STYLE NAMES IN REFINEMENT PATTERNS

denotes a required change to the associated abstract-level object. The intended meaning is obvious from context. A reference to a style in a refinement pattern is abbreviated according to the naming conventions summarized in Table I.

We assume that connections in an architecture do not share interface points. Multiple uses of a given interface point are modeled with multiple copies of the same point. This model has the advantage that interfaces and connections can be refined more flexibly. However, this choice of representation can result in an increase in the number of interface points.

A. Component Refinement

Figure 5 contains a refinement pattern for decomposing a functional component into a collection of components wrapped by a module. Component f is refined into module f' , hence the association $f \rightarrow f'$. A module signature contains all externally visible interfaces within the module. Since each interface point is an object with a unique name, there is no confusion as to the correspondences between the interface points of f and those of components in f' . By requiring that f and f' have the same signature, we are guaranteed that the original connections involving f are maintained through its subcomponents. The refinement is faithful because the interface requirement on f and f' prevents the addition or deletion of connections.

The next two patterns are for aggregating variables in situations that are common in intermediate stages of a development. This is done for time and space efficiency, es-

PATTERN OF ABSTRACT ARCHITECTURE:
$M: \text{MODULE}[p_1 \rightarrow p_2]$
COMPONENTS
$f: \text{FUNCTION}[p_{11} \rightarrow p_{12}]$
PATTERN OF CONCRETE ARCHITECTURE:
$M: \text{MODULE}[p_1 \rightarrow p_2]$
COMPONENTS
$f': \text{MODULE}[p_{11} \rightarrow p_{12}]$
ABSTRACT TO CONCRETE ASSOCIATIONS:
$f \rightarrow f'$

Fig. 5. Decomposing a component into subcomponents (Pattern 1)

pecially if the variables hold large objects. Application of the patterns also results in a simpler design.

Figure 6 contains a pattern for merging shared variables when one of them is a private variable. This pattern merges a shared variable m_1 , which is written by component f_1 and read by component f_2 , with a private variable m_2 , which is read and written by component f_1 . This is expressed by the association $(m_1, m_2) \rightarrow m'$. There are three basic requirements on this form of refinement:

- The variables denoted by schema variables m_1 and m_2 must have the same type, denoted by schema variable t .
- Only the component denoted by f_1 can write the variable denoted by m_1 . This prevents a new flow to f_1 , which would violate the faithfulness requirement.
- Only f_1 accesses private variable m_2 , otherwise a new flow would be created by the refinement. This requirement is enforced by the constraint on the abstract architecture.

A variant of this pattern combines the shared variable and the private variable into two fields of a record structure. With this variant, the constraint on the abstract architecture is not needed, provided that the components involved access only the proper fields of the record. This kind of refinement would not increase efficiency, but could help simplify the design.

Figure 7 contains a pattern for merging shared variables when neither of them are private. The two shared variables are connected by a common functional component. A shared variable denoted by schema variable m_1 is written by functional component f_1 and read by f_2 . Shared variable m_2 is written by f_2 and read by f_3 . The merge is expressed by the association $(m_1, m_2) \rightarrow m'$.

Our correctness criterion places the following restrictions on the architectures:

- The variables to be merged must be of the same type t .
- Since we treat dataflow as an intransitive relation, we also treat other relations dealing with the flow of data as intransitive relations. Therefore, functional components f_1 , f_2 , and f_3 have to be executed sequentially in batch mode so that we cannot infer the existence of a new abstract flow from f_1 to f_3 . This is prevented by configuration assertions a'_5 and a'_6 .
- No other functional components can read m_1 or write




Fig. 6. Merging a shared variable with a private variable (Pattern 2)

m_2 , which is enforced by a constraint on the abstract architecture.

A variant of this pattern combines the shared variables into two fields of a record structure. With this variant, the sequential ordering assertions in the concrete architecture and the constraint on the abstract architecture are not needed.

B. Connector Refinement

Figure 8 contains a pattern for implementing a dataflow connector by a pipe. Dataflow channel c from f_1 to f_2 is refined into a pipe c' connecting f_1 to f_2 . The connector refinement is expressed by the associations $c \dashrightarrow c'$ and $a \dashrightarrow a'$. This refinement is obviously faithful. Semantically, it can be justified on the basis of the meaning of the dataflow and pipe connectors.

Figure 9 contains a pattern for refining two functional components f_1 and f_2 that are executed in batch-sequential mode into a module with a main functional component f' transferring control first to f_1 and then to f_2 . The correctness of refinements of this form depends on the following properties.

- Component f_1 has to complete before f_2 can start, which is enforced by configuration assertion a' .
- Concrete component f' cannot transfer control to f_2 until f_1 completes, and f_1 cannot transfer control to f' after f_2 starts. These ordering relationships are




Fig. 7. Merging shared variables (Pattern 3)




Fig. 8. Implementing a dataflow connector by a pipe (Pattern 4)

PATTERN OF ABSTRACT ARCHITECTURE:
M: MODULE[$p_1 \rightarrow p_2$]
COMPONENTS
$f_1: F!Function[p_{11} \rightarrow p_{12}]$
$f_2: F!Function[p_{21} \rightarrow p_{22}]$
CONFIGURATION
$a: BS!Starts_After_Finish_Of(f_1, f_2)$
PATTERN OF CONCRETE ARCHITECTURE:
M: MODULE[$p_1 \rightarrow p_2$]
COMPONENTS
$f': F!Function[- \rightarrow -]$
$f_1: F!Function[p_{11} \rightarrow p_{12}]$
$f_2: F!Function[p_{21} \rightarrow p_{22}]$
CONNECTORS
$s'_1: CT!Enabling_Signal$
$s'_2: CT!Enabling_Signal$
CONFIGURATION
$a'_1: CT!Sender(s'_1, f_1)$
$a'_2: CT!Receiver_Signal(s'_1, f')$
$a'_3: CT!Sender(s'_2, f_2)$
$a'_4: CT!Receiver_Signal(s'_2, f_2)$
$a': CT!Before(s'_1, s'_2)$
ABSTRACT TO CONCRETE ASSOCIATIONS:
$a \dashrightarrow a'$
CONSTRAINTS ON CONCRETE ARCHITECTURE:
$\neg(\exists s': CT!Enabling_Signal [CT!Sender(s', f') \wedge CT!Receiver_Signal(s', f_2) \wedge CT!Before(s', s'_1)])$
$\neg(\exists s': CT!Enabling_Signal [CT!Sender(s', f_1) \wedge CT!Receiver_Signal(s', f') \wedge CT!Before(s'_2, s')])$

Fig. 9. Implementing ordering constraint using explicit control transfer (Pattern 5)

enforced by the two constraints on the concrete architecture.

- All functional components have to be enabled by f' and every control transfer must be between f' and a functional component. This is enforced by a well-formedness constraint in the control-transfer style, not by a constraint in the pattern.

C. Interface Refinement

Figure 10 contains the full specification of the pattern introduced earlier in Figure 3. The refinement of the dataflow connection into a shared-memory implementation has the side effect of changing the signature of the two functions, since connections do not share interface points.

VIII. EXAMPLE REVISITED

We now apply the refinement patterns to the compiler architectures illustrated earlier in Figure 1. In particular, we show how the level-1 compiler architecture can be refined into the level-2 compiler architecture using five of the patterns. The textual specification of the architectures are simplified through the use of ellipses for parts of the specification that are not relevant to the refinement under consideration. The full textual specifications for levels 1 and 2 are in Figure 2 and the appendix, respectively.

PATTERN OF ABSTRACT ARCHITECTURE:
M: MODULE[$ip:t, p_1 \rightarrow op:t, p_2$]
COMPONENTS
$f_1: F!Function[p_{11} \rightarrow op:t, p_{12}]$
$f_2: F!Function[ip:t, p_{21} \rightarrow p_{22}]$
CONNECTORS
$c: D!Dataflow_Channel[t]$
CONFIGURATION
$a: D!Connects(c, op, ip)$
PATTERN OF CONCRETE ARCHITECTURE:
M: MODULE[$p_1 \rightarrow p_2$]
COMPONENTS
$f_1: F!Function[p_{11} \rightarrow p_{12}]$
$f_2: F!Function[p_{21} \rightarrow p_{22}]$
$m': SM!Variable[t]$
CONFIGURATION
$a'_1: SM!Writes(f_1, m')$
$a'_2: SM!Reads(f_2, m')$
ABSTRACT TO CONCRETE ASSOCIATIONS:
$c \dashrightarrow m' \quad a \dashrightarrow (a'_1, a'_2)$
$(op, ip) \dashrightarrow$

Fig. 10. Implementing dataflow with a shared variable (Pattern 6)

The development of the level-2 architecture involves three main steps — the introduction of the pipe between the lexical analyzer and the parser, the development of the shared tree accessed by the parser, analyzer/optimizer, and code generator, and the development of the shared symbol table between the lexical analyzer and the optimizer. All patterns, with the exception of Pattern 5, are used. (Pattern 5 is applied repeatedly to the level-2 compiler architecture to get the level-3 architecture in the appendix.)

A. Introduction of the Pipe

This refinement is a straightforward application of Pattern 4. Consider the following abbreviated subarchitecture of the level-1 compiler.

```
compiler_L1: MODULE
    [char_iport: SEQ(character) -> code_oport: code]
    COMPONENTS
        lexical_analyzer: Function
        [... -> token_oport: SEQ(token), ...]
        parser: Function[token_iport: SEQ(token) -> ...]
    CONNECTORS
        token_channel: Dataflow_Channel[SEQ(token)]
    CONFIGURATION
        token_flow:
            Connects(token_channel, token_oport, token_iport)
```

Pattern 4 can be used to refine dataflow channel `token_channel` into pipe `token_pipe`, resulting in the following level-2 architecture.⁹

```
compiler_L2: MODULE
    [char_iport: SEQ(character) -> code_oport: code]
    COMPONENTS
        lexical_analyzer_module: MODULE
        [... -> token_oport: Finite_Stream(token)]
        parser: Function
        [token_iport: Finite_Stream(token) -> ]
```

⁹An output and an input port of type `SEQ(token)` were implemented as type `Finite_Stream(token)`. A stream is a function from clock times to values. The correctness of this type refinement is not treated in this paper.

```

CONNECTORS
  token_pipe: Pipe[Finite_Stream(token)]
CONFIGURATION
  token_flow:
    Connects(token_pipe, token_oport, token_iport)

```

B. Development of the Shared Abstract Syntax Tree

Consider the following dataflow architecture.

```

compiler_L1: MODULE
  [char_iport: SEQ(character) -> code_oport: code]
COMPONENTS
  parser: Function [... -> base_ast_oport: ast]
  analyzer_optimizer: Function
    [base_ast_iport: ast, ... -> full_ast_oport: ast]
  code_generator: Function
    [full_ast_iport: ast -> ...]
CONNECTORS
  base_ast_channel: Dataflow_Channel[ast]
  full_ast_channel: Dataflow_Channel[ast]
CONFIGURATION
  base_ast_flow:
    Connects(base_ast_channel,
              base_ast_oport, base_ast_iport)
  full_ast_flow:
    Connects(full_ast_channel,
              full_ast_oport, full_ast_iport)

```

It can be split into two dataflow architectures and Pattern 6 is applied to each to construct two shared memory architectures, which are composed horizontally to form a single architecture. Then, Pattern 3 can be applied to merge the two shared data structures into a single shared tree, called `abstract_syntax_tree`. The three architectures compose vertically, so we know that the final architecture, given below, is correct with respect to the original dataflow architecture.

```

compiler_L2: MODULE
  [char_iport: SEQ(character) -> code_oport: code]
COMPONENTS
  parser: Function[... -> ]
  analyzer_optimizer: Function[ -> ]
  code_generator: Function[ -> ...]
  abstract_syntax_tree: Variable[ast]
CONFIGURATION
  write_base_ast:
    Writes(parser, abstract_syntax_tree)
  read_base_ast:
    Reads(analyzer_optimizer, abstract_syntax_tree)
  write_full_ast:
    Writes(analyzer_optimizer, abstract_syntax_tree)
  read_full_ast:
    Reads(code_generator, abstract_syntax_tree)
precedence_1:
  Starts_After_Finish_Of(analyzer_optimizer, parser)
precedence_2:
  Starts_After_Finish_Of(code_generator,
                        analyzer_optimizer)

```

C. Development of the Shared Symbol Table

This refinement involves three individual refinements, but only vertical composition. Consider the following architecture, which specifies the dataflow from the lexical analyzer to the analyzer/optimizer that is used to transmit binding information.

```

compiler_L1: MODULE
  [char_iport: SEQ(character) -> code_oport: code]
COMPONENTS

```

```

lexical_analyzer: Function
  [char_iport: SEQ(character)
   -> bind_oport: SEQ(binding), ...]
  analyzer_optimizer: Function
    [..., bind_iport: SEQ(binding) -> ...]
CONNECTORS
  bind_channel: Dataflow_Channel[SEQ(binding)]
CONFIGURATION
  bind_flow:
    Connects(bind_channel, bind_oport, bind_iport)

```

The three refinement steps are:

1. Pattern 1 is used to refine the lexical analyzer into a new module containing itself and a private symbol table used to store bindings locally before proceeding to the next phases of compilation, which could modify the table.
2. Pattern 6 is used to introduce a shared variable between the lexical analyzer and the optimizer, corresponding to `bind_channel`, that can be used to transmit the completed symbol table.¹⁰
3. Pattern 2 is used to merge the private symbol table and the shared variable into a single shared repository. This reflects a conscious decision to allow no component other than the lexical analyzer to write the table. As a consequence, any additional information, such as storage requirements, and code restructuring must be represented in the abstract syntax tree.

The resulting architecture is given below.

```

compiler_L2: MODULE
  [char_iport: SEQ(character) -> code_oport: code]
COMPONENTS
  lexical_analyzer_module: MODULE[... -> ...]
  COMPONENTS
    lexical_analyzer: Function[... -> ...]
    symbol_table: Variable[SEQ(binding)]
  CONFIGURATION
    write_bind:
      Writes(lexical_analyzer, symbol_table)
    read_bind:
      Reads(lexical_analyzer, symbol_table)
  END lexical_analyzer_module
  analyzer_optimizer: Function[ -> ]
CONFIGURATION
  read_bind:
    Reads(analyzer_optimizer,
          lexical_analyzer_module!symbol_table)

```

D. Putting The Pieces Together

The three individual architecture hierarchies can be flattened to two levels because faithful interpretations are transitive. Then, they can be composed horizontally to form the composite compiler architectures at levels 1 and 2. The level-3 compiler architecture can be formed in a similar fashion.

It is worth noting that a series of refinements can result in a deep hierarchy that need not be saved explicitly. The sequence of steps in deriving a concrete architecture are important, but the intermediate architectures themselves may not be. We saw this in the development of the symbol table.

¹⁰The nested `lexical_analyzer_module` can be flattened by a restructuring pattern so that patterns can be applied directly.

We also observe that it is possible to adopt a hybrid approach to architecture development in which parts of the architecture are developed by means of refinements and other parts are specified completely by hand. In the latter situation, refinement patterns can be used to validate the correctness of the putative implementation architectures through a straightforward matching procedure. Correct hierarchies can be composed no matter how they were developed, provided the composition is faithful.

IX. APPLICATION TO A POWER-CONTROL SYSTEM

The approach presented in this paper has been used to design an architecture for an operational power control system implemented in 200,000 lines of FORTRAN 77 code. The system is used by Tokyo Electric Power Company, Inc. to achieve efficient administration of power-supply systems in Tokyo, Japan. The power-control system was developed by Meidensha Corporation and its architecture is considered a company asset. Originally, the details of the architecture were represented informally in several loosely connected documents. This created a difficult situation for Meidensha Corp. because they wanted to expand their business in control systems to other areas with similar requirements, which would require minor modifications to the reference architecture. With no formalized architecture, such an expansion would certainly lead to duplication of effort and unnecessary errors in implementation.

Our objective was to formalize the reference architecture in terms of company styles and at two levels of detail, and to guarantee that the concrete architecture is correct with respect to the abstract architecture. This task was completed successfully. The abstract architecture was stated in terms of a dataflow style, and the concrete architecture was a combination of a call-return style, a (structured) shared-memory style, and a special process synchronization style for DEC VMS operating systems. Twelve patterns were used in the development; each was used many times.

Pattern 1 was used for decomposing functional components into modules; Pattern 6 was used to implement dataflow as a shared variable. Domain-specific refinement patterns were needed to handle two distinctive features of the concrete power-control architecture—heavy use of shared memory and process synchronization by an enabling signal. The shared memory did not have a uniform structure. Dozens of dataflows were implemented by a single record containing one field for each flow. Some dataflows were implemented as a record structure containing the data and a one-bit enabling signal, and others as a message channel plus a signaling channel. A collection of variables containing one bit are packaged into a bitstring for efficient communication. Variants of Patterns 2 and 3 were used to aggregate individual variables into records.

This successful experience strongly suggests that, in the domain of power control, only a small number of patterns is required. This allows the cost of pattern verification to be amortized across many applications in the power-control domain. We know that many of the patterns are relevant in other domains as well, and believe that only a modest

number of new patterns will be needed in many application areas.

X. RELATED WORK

The field of architecture-driven software development will not reach its full potential until it is possible to refine and compose architectures incrementally, flexibly, and in ways that preserve the desired properties. Ideally, deep properties of an architecture, such as relative correctness, should be preserved. This requires that an architecture hierarchy be represented formally and the mapping between the levels be precise and explicit. We review related work in the areas of refinement, correctness, and composition.

Previous approaches to specification refinement have concentrated on the preservation of functional properties, which occurs when the mapping between specifications is a theory interpretation. The mapping often is complicated by a change in data representation. This can be taken into account by adapting the technique of Hoare [12] to relate the types in the abstract and concrete specifications. An analogous problem arises in architecture refinement when there is a change in style. We have introduced the notion of a style mapping to relate the styles in the abstract and concrete architectures.

We are not the first to recognize the importance of schematic transformations in stepwise refinement. In [10], Gerhart gives several examples of schema transformations that preserve functional correctness. We define schema transformations that preserve architecture correctness. The two forms of refinement are complementary. An architecture refinement hierarchy describes system organization — its components, interfaces, and connections. Functional refinement is used to develop the behavior of the system components in the architecture. In both instances, schemas can be used to increase the reusability of designs and proofs.

Of course, the utility of architecture hierarchies has been recognized for some time. For example, in the 1970s Jackson [13], Yourdan and Constantine [20], DeMarco [7], and others describe system architectures and, more recently, architectural description has been the basis for commercial offerings. However, previous work has given little attention to the mapping between levels of abstraction. We formally defined the interpretation mapping required in architecture correctness proofs in terms of a specific name mapping and a general, reusable style mapping. The mapping also provides the basis for traceability of architectural design decisions, which is useful in practice.

Recently, another form of a mapping between architectures has been developed for the Rapide architecture definition language [14], [15]. Rapide is used to define executable architectures based on distributed event processing. Two architectures are related by mapping concrete events to abstract events. Event mappings provide the basis for comparative simulation, a technique that complements static modeling.

The standard criterion for functional correctness is not applicable to architectures because of the completeness as-

sumption. A similar completeness assumption is made widely in the database community for analogous reasons, see Reiter [19]. However, Reiter allows only finitely many objects, so a “domain closure axiom” can be used to enumerate the domain of discourse. No similar technique can be applied here because, in general, an architecture can be infinite. For example, we allow quantification over infinite types (such as integers) and dynamic architectures with an unbounded number of processes and connections. Because of the completeness assumption, an abstract architecture must be faithfully interpreted in the concrete architecture.

In [17], Moriconi and Hare study the relative correctness of two architectures under the completeness assumption. They make the simplifying assumption that an architecture can contain only a fixed, finite number of objects. Broy [5], Brinksma [4], and others have applied the standard approach to correctness to architectures. Broy’s component refinements turn out to be conservative (and, hence, faithful) because interface signatures are preserved, but his connection refinements may not be because additional flows could be added to a channel. Brinksma justifies channel splitting on the basis of behavioral reasoning; application of his rule can violate the completeness assumption.

We appear to be the first to observe that, in an architectural correctness proof, it is important to establish the semantic correctness of the relevant style mappings. The importance of reasoning about connectors was recognized by Allen and Garlan [3], who formalize them in a subset of CSP [11] and then proved absence of deadlock. In [18] we define the meaning of connectors axiomatically in a temporal logic and prove both fairness and safety properties of an implementation of the dataflow connector in shared memory. Garlan et al [8], [9] also have done important work on identifying and exploiting architectural styles. We build on their work, developing schematic style mappings and schematic refinements involving style-to-style transformations.

Composition has been studied recently by Abadi and Lamport [1], [2]. Their results are semantic and applicable to any domain, whereas ours are syntactic and specialized to the domain of software architecture. It is easy to state general criteria for the correctness of horizontal composition of architectures. However, it requires a difficult proof that it is not possible to infer new facts about the composite abstract architecture from the composite concrete architecture. Therefore, we defined a new specialized form of horizontal composition that requires only very simple syntactic checks. Broy [5] gives three operators for composing functional-style architectures, but does not consider the composition of architectures involving multiple styles. Vertical composition in a hierarchy of architectures is immediate provided each level in the hierarchy is correct with respect to the immediately preceding level.

XI. CONCLUSION

We have described a stepwise refinement methodology for the development of a heterogeneous hierarchy of architectures that are relatively correct under a particular

completeness assumption. We introduced the notion of an architecture refinement pattern as the principal vehicle for codifying reusable solutions to routine architectural design problems. Once an architecture refinement pattern is proved correct, instances of it can be used in a particular development with no further proof. Patterns are compositional and can be proved in isolation. Subsystem architectures are compositional provided they overlap only in certain ways. The methodology was used successfully to explicate the architectural design of an operational power-control system.

To develop a theory of correctness for architecture refinement, we adapted the technique of faithful interpretation that was introduced in an earlier paper for after-the-fact verification of complete architectures [18]. A new proof technique for checking faithfulness was presented. The interpretation mapping between architectures was simplified by decomposing it into an architecture-specific name mapping and a general style-to-style mapping. We are not aware of this distinction being made elsewhere in the literature. It is important because a style mapping and its proof, both of which can be complex, can be reused in validating any pattern involving the two styles. In contrast, a name mapping is simple, specific to a pattern, and cannot be validated independent of the pattern.

An important premise behind our work is that at least the dominant styles of architectural design can be generalized to partially interpreted schema and most architecture refinements for these styles can be generalized to transformations on schema. We believe that a small number of architectural styles are sufficient for a large number of application domains, and that only a modest number of refinement patterns are needed between each pair of styles. This assertion is supported to some degree by the experiences reported in this paper regarding the compiler and power-control architectures.

Some methodological implications of our faithfulness requirement are worth mentioning. First, architectural styles should clearly differentiate among different architectural concepts. Consider a transaction on a distributed database system, which is an atomic operation logically but rarely is a physically atomic operation. If the abstract “transaction” connector is refined into a two-phase commit protocol involving a series of data transmissions, the refinement will not be faithful unless the purpose of the two-phase commit is taken into account in the design of the style. For example, the commit protocol can be modeled in terms of special “control” connectors that are distinct from the connector that models the transfer of data from the database to the designated site. Then, the abstract flow of data will be the same as the concrete flow, even though there is extra preparatory activity in the concrete architecture. Second, architects can, but should not, circumvent the completeness assumption by adding concepts to a concrete architecture that are unrelated to those in the associated abstract architecture. A correctness criterion could be defined that disallows this, but it would be too restrictive for both design and composition. It is the sort of thing that is unlikely

to happen by accident. However, the only real safeguard is the careful scrutiny of each refinement pattern.

We have completed an initial implementation of our methodology sufficient to demonstrate its feasibility. The tool accepts as input a collection of refinement patterns, an abstract architecture, and a concrete architecture. The tool matches instances of the patterns on the abstract and concrete architectures with no user intervention. It makes no attempt to generate instances at this time. One correct composition of refinements is found, if it exists, although in general there may be many possible correct compositions. Specific failures are reported if there is not complete coverage. Any constraints on the application of a refinement pattern are checked automatically. This tool was used in the compiler and the power-control application.

Future work involves the development and evaluation of a handbook of architectural refinement patterns. Good designers tend to use well-established architectural styles, including both basic idioms (such as pipe-filter, client-server, and layering) and reference models (such as the ISO OSI 7-layer model [16]). We are now expanding our library to relate more styles as well as to elaborate more configurations involving the styles in the paper. Eventually, we would like to have a large enough library to support "industrial strength" architecture design. For example, we would like to be able to start with an abstract architecture for a large system, in say a dataflow style, refine it into architectures in a dominant commercial style, such as client/server, and then refine that architecture into an implementation-level architecture that specifies the exact forms of communication. In developing a pattern library, we will be concerned with more than correctness. In particular, we want to use architectural refinement patterns to achieve a greater degree of system predictability. For example, it would be useful to have refinement patterns that optimize performance for specific processors or, more generally, for a given computing and network environment.

Our longer-term objective is to develop a practical architecture synthesis tool that is driven by a broadly useful pattern library. The tool will enforce a design discipline similar to the one enforced by commercial hardware synthesis tools. These tools gain much of their power from the use of clearly defined and reusable styles: typically, register-transfer, logic, and gate-level styles. A pattern library of the sort proposed in this paper is expected to enable effective synthesis of software architectures.

APPENDIX

I. LOWER LEVEL COMPILER ARCHITECTURES

The textual specifications for the two implementations of the compiler architecture make extensive use of imported types and styles, which are not defined in this paper. The specifications have a straightforward translation into logic. The following is the full level-2 specification.

```
compiler_L2: MODULE
  [char_iport: SEQ(character) -> code_oport: code]
  IMPORT character, code, token, binding, ast
    FROM compiler_types
```

```
  IMPORT Function FROM Functional_Style
  IMPORT Pipe, Finite_Stream, Connects
    FROM Process_Pipeline_Style
  IMPORT Variable, Reads, Writes
    FROM Shared_Memory_Style
  IMPORT Start_After_Finish_Of
    FROM Batch_Sequential_Style
COMPONENTS
  lexical_analyzer_module: MODULE
    [char_iport: SEQ(character)
      -> token_oport: Finite_Stream(token)]
    EXPORTING lexical_analyzer, symbol_table
      IMPORT character, token, binding
        FROM compiler_types
    IMPORT Function FROM Functional_Style
    IMPORT Variable, Reads, Writes
      FROM Shared_Memory_Style
COMPONENTS
  lexical_analyzer: Function
    [char_iport: SEQ(character)
      -> token_oport: Finite_Stream(token)]
    symbol_table: Variable[SEQ(binding)]
CONFIGURATION
  write_bind:
    Writes(lexical_analyzer, symbol_table)
  read_bind:
    Reads(lexical_analyzer, symbol_table)
END lexical_analyzer_module
parser:
  Function[token_iport: Finite_Stream(token) -> ]
  analyzer_optimizer: Function[ -> ]
  code_generator: Function[ -> code_oport: code]
  abstract_syntax_tree: Variable[ast]
CONNECTORS
  token_pipe: Pipe[Finite_Stream(token)]
CONFIGURATION
  token_flow:
    Connects(token_pipe, token_oport, token_iport)
  read_bind:
    Reads(analyzer_optimizer,
      lexical_analyzer_module!symbol_table)
  write_base_ast: Writes(parser, abstract_syntax_tree)
  read_base_ast:
    Reads(analyzer_optimizer, abstract_syntax_tree)
  write_full_ast:
    Writes(analyzer_optimizer, abstract_syntax_tree)
  read_full_ast:
    Reads(code_generator, abstract_syntax_tree)
  precedence_1:
    Starts_After_Finish_Of(analyzer_optimizer, parser)
  precedence_2:
    Starts_After_Finish_Of(code_generator,
      analyzer_optimizer)
END compiler_L2
```

The level-3 compiler architecture employs a common implementation of the batch-sequential style. In particular, the batch processing in the level-2 compiler is implemented in terms of a main program and subroutines, as illustrated in Figure 11. This implementation is justified by Pattern 5, which was presented in the body of the paper.

The wiring at level 3 is constrained by the temporal-precedence assertions at level 2.

```
precedence_1:
  Starts_After_Finish_Of(analyzer_optimizer, parser)
precedence_2:
  Starts_After_Finish_Of(code_generator,
    analyzer_optimizer)
```

We have to make sure that the transfer of control satisfies this temporal ordering of the computation. Two ap-




Fig. 11. Third level in architecture hierarchy for compiler

plications of Pattern 5 can be used to guarantee that the ordering relations are satisfied independently. The horizontal composition of the two applications of Pattern 5 guarantees that the composite architecture satisfies both orderings.

The composite level-3 architecture is given below.

```
compiler_L3: MODULE
  [char_iport: SEQ(character) -> code_oport: code]
  IMPORT character, code, token, binding, ast
  FROM compiler_types
  IMPORT Function FROM Functional_Style
  IMPORT Pipe, Finite_Stream, Connects
  FROM Process_Pipeline_Style
  IMPORT Variable, Reads, Writes
  FROM Shared_Memory_Style
  IMPORT Enabling_Signal, Sender, Receiver, Before
  FROM Control_Transfer_Style
COMPONENTS
  main: Function[ -> ]
  lexical_analyzer_module: MODULE
    [char_iport: SEQ(character)
     -> token_oport: Finite_Stream(token)]
    EXPORTING lexical_analyzer, symbol_table
    IMPORT character, token, binding
    FROM compiler_types
    IMPORT Function FROM Functional_Style
    IMPORT Variable, Reads, Writes
    FROM Shared_Memory_Style
COMPONENTS
  lexical_analyzer: Function
    [char_iport: SEQ(character)
     -> token_oport: Finite_Stream(token)]
    symbol_table: Variable[SEQ(binding)]
CONFIGURATION
  write_bind:
    Writes(lexical_analyzer, symbol_table)
  read_bind:
    Reads(lexical_analyzer, symbol_table)
END lexical_analyzer_module
parser:
  Function[token_iport: Finite_Stream(token) -> ]
  analyzer_optimizer: Function[ -> ]
  code_generator:
    Function[ -> code_oport: code]
    abstract_syntax_tree: Variable[ast]
CONNECTORS
  token_pipe: Pipe[Finite_Stream(token)]
  start_main, start_lex, start_parse, parse_finish,
  start_opt, opt_finish, start_gen, gen_finish,
  main_finish: Enabling_Signal
CONFIGURATION
  token_flow:
    Connects(token_pipe, token_oport, token_iport)
  read_bind:
```

```
  Reads(analyzer_optimizer,
        lexical_analyzer_module!symbol_table)
  write_base_ast:
    Writes(parser, abstract_syntax_tree)
  read_base_ast:
    Reads(analyzer_optimizer, abstract_syntax_tree)
  write_full_ast:
    Writes(analyzer_optimizer, abstract_syntax_tree)
  read_full_ast:
    Reads(code_generator, abstract_syntax_tree)
  rcvr_start_main: Receiver(start_main, main)
  sndr_start_lex: Sender(start_lex, main)
  rcvr_start_lex:
    Receiver(start_lex,
            lexical_analyzer_module!lexical_analyzer)
  sndr_start_parse: Sender(start_parse, main)
  rcvr_start_parse: Receiver(start_parse, parser)
  sndr_parse_finish: Sender(parse_finish, parser)
  rcvr_parse_finish: Receiver(parse_finish, main)
  sndr_start_opt: Sender(start_opt, main)
  rcvr_start_opt:
    Receiver(start_opt, analyzer_optimizer)
  sndr_opt_finish:
    Sender(opt_finish, analyzer_optimizer)
  rcvr_opt_finish: Receiver(opt_finish, main)
  sndr_start_gen: Sender(start_gen, main)
  rcvr_start_gen:
    Receiver(start_gen, code_generator)
  sndr_gen_finish: Sender(gen_finish, code_generator)
  rcvr_gen_finish: Receiver(gen_finish, main)
  sndr_main_finish: Sender(main_finish, main)

  start_main_before_lex:
    Before(start_main, start_lex)
  start_main_before_parse:
    Before(start_main, start_parse)
  start_parse_before_finish:
    Before(start_parse, parse_finish)
  finish_parse_before_start_opt:
    Before(parse_finish, start_opt)
  start_opt_before_finish:
    Before(start_opt, opt_finish)
  finish_opt_before_start_gen:
    Before(opt_finish, start_gen)
  start_gen_before_finish:
    Before(start_gen, gen_finish)
  finish_gen_before_main:
    Before(gen_finish, main_finish)
END compiler_L3
```


The associations between these two levels are

```
precedence_1 --> finish_parse_before_start_opt
precedence_2 --> finish_opt_before_start_gen
```

REFERENCES


- [1] M. Abadi and L. Lamport, "Composing Specifications", *ACM Transactions on Programming Languages and Systems*, Vol. 15, No. 1, January 1993, pp. 73-132.
- [2] M. Abadi and L. Lamport, "Conjoining Specifications", Technical Report 118, Digital Systems Research Center, Palo Alto, California, December 1993.
- [3] R. Allen and D. Garlan, "Formalizing Architectural Connection", *Proceedings of the Sixteenth International Conference on Software Engineering*, May 1994, pp. 71-80.
- [4] E. Brinksma, B. Jonsson, and F. Orava, "Refining Interfaces of Communicating Systems", *TAPSOFT'91: Lecture Notes in Computer Science 494*, S. Abramsky and T.S.E. Maibaum, Eds., Springer-Verlag, 1991, pp. 297-312.
- [5] M. Broy, "Compositional Refinement of Interactive Systems", No. 89, Digital Systems Research Center, Palo Alto, California, July 1992.
- [6] H. B. Enderton, *A Mathematical Introduction to Logic*, Academic Press, 1972.

- [7] T. DeMarco, *Structured Analysis and System Specification*, Yourdan Press, 1979.
- [8] D. Garlan, R. Allen, and J. Ockerbloom, "Exploiting Style in Architectural Design Environments", *Proceedings of ACM SIGSOFT'94: Symposium on Foundations of Software Engineering*, New Orleans, Louisiana, December 1994.
- [9] D. Garlan and M. Shaw, "An Introduction to Software Architecture", In *Advances in Software Engineering and Knowledge Engineering*, Volume 1, V. Ambriola and G. Tortora, Eds., World Scientific Publishing Company, 1993.
- [10] S.L. Gerhart, "Knowledge about programs", *Proceedings of the International Conference on Software Reliability*, Los Angeles, California, April 1975, pp. 88-95.
- [11] C.A.R. Hoare, *Communicating Sequential Processes*, Prentice-Hall, 1985.
- [12] C.A.R. Hoare, "Proof of correctness of data representations", *Acta Informatica*, Vol. 1, No. 4, 1972, pp. 271-281.
- [13] M.A. Jackson, *Principles of Program Design*, Academic Press, 1975.
- [14] D. Katiyar, D.C. Luckham, and J. Mitchell, "A type system for prototyping languages", *Proceedings of the 21st ACM Symposium on Principles of Programming Languages*, Portland, Oregon, 1994.
- [15] D.C. Luckham, J. Vera, D. Bryan, L. Augustin, and F. Belz, "Partial Orderings of Event Sets and Their Application to Prototyping Concurrent, Timed Systems", *Journal of Systems and Software*, Vol. 21, No. 3, June 1993, pp. 253-265.
- [16] G.R. McClain, editor, *Open Systems Interconnection Handbook*, McGraw-Hill, New York, N.Y., 1991.
- [17] M. Moriconi and D.F. Hare, "The PegaSys System: Pictures as Formal Documentation of Large Programs", *ACM Transactions on Programming Languages and Systems*, Vol. 8, No. 4, October 1986, pp. 524-546.
- [18] M. Moriconi and X. Qian, "Correctness and Composition of Software Architectures", *Proceedings of ACM SIGSOFT'94: Symposium on Foundations of Software Engineering*, New Orleans, Louisiana, December 1994.
- [19] R. Reiter, "Deductive Question-Answering on Relational Databases", in *Logic and Data Bases*, H. Gallaire and J. Minker, Eds., Plenum Press, 1978, pp. 149-177.
- [20] E. Yourdan and L.L. Constantine, *Structured Design: Fundamentals of a Discipline of Computer Program and Systems Design*, Prentice-Hall, Inc., Englewood Cliffs, N.J., 1979.



Xiaolei Qian received the B.Sc. degree from Xian Jiao Tong University, Xian, China, in 1982, and the M.Sc. and Ph.D. degrees from Stanford University, Stanford, CA, in 1984 and 1989, respectively, all in computer science.


She has been a Computer Scientist in the Computer Science Laboratory at SRI International since 1991. Before joining SRI, she was a Member of the Technical Staff at AT&T Bell Laboratories, and a Computer Scientist at Kestrel Institute. Her research interests include software architectures, semantic interoperability and integration of heterogeneous databases, and database security. She is also interested in database programming languages and formal methods.



R. A. Riemschneider received the B.S. (*summa cum laude*) degree in physics and mathematics from Miami University in 1973 and the M.A. degree in mathematics from the University of California at Berkeley in 1975.

He joined the Computer Science Laboratory of SRI International in 1991 as a Senior Software Engineer, where he performs research and development on applications of logic to software engineering. Prior to joining SRI, he was a Senior Research Scientist at Advanced Decision Systems, a founder of Reasoning Systems, a Computer Scientist at Systems Control Technology, and an Instructor at the University of California at Berkeley and the California State University at Hayward.

Mr. Riemschneider is a member of the Association for Symbolic Logic, the Association for Computing Machinery, and the IEEE Computer Society.



Mark Moriconi received a Ph.D. degree in computer science from the University of Texas at Austin in 1978.

He joined the Computer Science Laboratory of SRI International in 1978 and has been its Director since 1989. Prior to joining SRI, he was a research scientist at the University of Texas at Austin and a research assistant at USC Information Science Institute. His main research interests are in the use of formal methods in software development. He is currently working on formal methods for architecture-based software composition.

Dr. Moriconi is a member of the Association for Computing Machinery, and the IEEE Computer Society. He is on the editorial board of *IEEE Transactions on Software Engineering* and has served on numerous technical program committees in the areas of software engineering and formal methods. He is General Chair for the upcoming ACM SIGSOFT '96 Symposium on Foundations of Software Engineering, which will have a special focus on software architecture.

C SRI Publications: Correctness and Composition of Software Architectures

Correctness and Composition of Software Architectures*

Mark Moriconi and Xiaolei Qian

Computer Science Laboratory
SRI International
Menlo Park, California 94025

ABSTRACT

The design of a large system typically involves the development of a hierarchy of different but related architectures. A criterion for the relative correctness of an architecture is presented, and conditions for architecture composition are defined which ensure that the correctness of a composite architecture follows from the correctness of its parts. Both the criterion and the composition requirements reflect special considerations from the domain of software architecture.

The main points are illustrated by means of familiar architectures for a compiler. A proof of the relative correctness of two different compiler architectures shows how to decompose a proof into generic properties, which are proved once for every pair of architectural styles, and instance-level properties, which must be proved for every architecture.

1 Introduction

The development of an architecture for a large system is a complicated task that can be made simpler by means of a stepwise development methodology. Ideally, an architect would use a hierarchical approach in which the composition of lower-level architectures is guaranteed to implement a higher-level architecture. The foundations for such an approach must include a method for proving that one architecture implements another architecture and a means of composing architectures so that the composite architecture is correct if all of its components are correct. We examine both problems in

This research was supported by the Advanced Research Projects Agency under Rome Laboratory contract F30602-93-C-0245.

this paper. We work at the logic level, independent of a particular architecture definition language. Thus, our results can be applied to a large class of such languages.

An architecture hierarchy is a sequence of two or more individual architectures that may differ with respect to the number and kind of objects and connections. For example, an abstract architecture containing functional components related by dataflow connections may be implemented in a concrete architecture in terms of procedures, control connections, and shared variables. An abstract architecture usually is smaller and easier to understand; a concrete architecture usually reflects more implementation concerns. A given architecture can be homogeneous (consisting of one style) or heterogeneous (consisting of multiple styles). Garlan and Shaw [7] provide a taxonomy of some common styles, including dataflow, pipe-and-filter, client-server, and event-based systems.

Before we can consider the relative correctness of two architectures, we first must decide on the meaning of the architectures. Suppose that, to facilitate system upgrades and maintenance on a particular system, we design a pipeline architecture that restricts the system topology to a linear sequence of filters. If a concrete architecture implements the pipeline, but additionally introduces feedback loops, the *raison d'être* behind the original pipeline architecture is no longer valid. In effect, there is no reason to specify a pipeline in the first place if all possible feedback loops are allowed in its implementation.

Therefore, we make a *completeness assumption* about a given architecture. Informally, the assumption is that, if an architectural fact is not explicit in the architecture, or deducible from the architecture, then the fact is not intended to be true of the architecture. In the pipeline example, it is not possible to infer the existence of a feedback loop from the linearity constraint, so we assume that no feedback loop is allowed in an implementation of the architecture. In general, an architecture (whether static or dynamic) can contain an unbounded number of facts.

The completeness assumption requires a correctness criterion that differs from the standard one (that is based on theory extension). In our application of the correctness criterion, we make a clear distinction between type-level properties that must be proved only once for every pair of architectural styles and instance-level properties that must be proved for every pair of architectures. This decomposition greatly simplifies correctness proofs and the statement of the mapping between two architectures. Composition is possible under the completeness assumption provided that certain syntactic constraints are satisfied.

This paper is organized as follows. The next two sections introduce basic architectural concepts and illustrate the correctness problem for architectures. Section 4 defines the correctness criterion in terms of logical theories, independent of any particular architectural definition language. Sections 5–7 explain how to use the criterion. Of particular interest is the construction and validation of the mapping between architectures. Section 8 defines necessary and sufficient conditions for architecture composition and defines two specific composition operators. Section 9 discusses related work, and the conclusion summarizes our results and discusses their possible implications for future research in software architecture.

2 Basic Architectural Concepts and Notation

A software architecture is represented using the following concepts.

1. **Component:** An object with independent existence, e.g., a module, process, procedure, or variable.
2. **Interface:** A typed object that is a logical point of interaction between a component and its environment.
3. **Connector:** A typed object relating interface points, components, or both.
4. **Configuration:** A collection of constraints that wire objects into a specific architecture.
5. **Mapping:** A relation between the vocabularies and the formulas of an abstract and a concrete architecture. The formula mapping is required because the two architectures can be written in different styles.
6. **Architectural style:** A style consists of a vocabulary of design elements, a set of well-formedness constraints that must be satisfied by any architecture written in the style, and a semantic interpretation of the connectors.

Components, interfaces, and connectors are treated as *first-class objects* — i.e., they have a name and they are refineable. Abstract architectural objects can be decomposed, aggregated, or eliminated in a concrete architecture. The semantics of components is not considered part of an architecture, but the semantics of connectors is.

We will use a simple notation for describing an architecture. Suppose that we want to describe the interaction between the parser and the semantic analyzer in a standard compiler. A dataflow architecture for this interaction is contained in Figure 1.¹

```

parse_analyze: MODULE
  IMPORT ...
  EXPORT ...
  COMPONENTS
    parser      : Function
    analyzer    : Function
  INTERFACES
    oast        : OPORT [ast] OF parser
    iast        : IPORT [ast] OF analyzer
  CONNECTORS
    ast_channel : Dataflow_Channel[ast]
  CONFIGURATION
    Connects(ast_channel, oast, iast)
END parse_analyze

```

Figure 1: Example Dataflow Architecture

The parser and analyzer are modeled as functional components. The parser (which accepts a sequence of tokens) has an output port *oast* that supplies an abstract syntax tree. The analyzer accepts a values of type *ast* (producing values of the same type). The dataflow connection is wired to the right ports by the assertion

Connects(ast_channel, oast, iast)

where *Connects(c,o,i)* means that connection *c* links output port *o* to input port *i*. All of the objects that make up the architecture are wrapped by a module, which can selectively import and export objects. In this example, we import some useful compiler types and the predefined functional and dataflow styles.

The dataflow architecture separates and names all components, ports, and connections. Observe that the signature of a component is not hard-wired to the component. A signature consists of individual ports that can be referenced and refined independently of the associated component. Interface separation will be useful later for architecture composition.

¹The precise syntax is not important for the purposes of this paper. Later, we formalize this architecture in logic, and that is the representation that is intended to express the intentions of the designer.

3 Illustration of the Problem

Suppose that we want to design the architecture for a compiler. A standard dataflow model of a compiler is depicted at the top of Figure 2. The diagram is used only as an informal pedagogical aid; it is not intended to be a formal specification. Boxes denote functional components and arrows denote directional dataflow between ports. The labels on arrows denote types or value domains. An object cannot be transmitted between ports unless its type is compatible with the types of the ports. The diagram is assumed to be complete in that there can be no other functional components, ports, or data flows.²




Figure 2: Two architectures for a compiler

Figure 2 also contains a concrete, hybrid architecture for the compiler that implements the dataflow style in terms of pipe-filter, batch-sequential, and shared-memory styles. Abstract signatures are changed in the concrete architecture, dataflow connections are implemented in several ways, through a pipe and shared data objects, and precedence relations are used to prevent direct flow of data from the parser to the code generator.

To illustrate the correctness problem, we focus on the implementation of the dataflow channel between the parser and analyzer in terms of the reading and writing of a shared abstract syntax tree. The implementation architecture is described textually in Figure 3. The shared abstract syntax tree is represented as a variable.³ The read and write relations are not named; they are primitives that cannot be refined.

²A dataflow connection is treated here as an intransitive relation.

³The shared abstract syntax tree might be represented as an encapsulated data type in a real compiler. If we had chosen that representation, the architecture would involve calls to access functions that read and write the internal variable used to represent the tree.

```
concrete_parse_analyze: MODULE
  IMPORT ...
  EXPORT ...
  COMPONENTS
    parser      : Function
    analyzer   : Function
    tree       : Variable[ast]
  CONFIGURATION
    Writes(parser, tree)
    Reads(analyzer, tree)
  END concrete_parse_analyze
```

Figure 3: Concrete Shared-Memory Architecture

The intended associations between the two architectures are

```
oast -->
iast -->
ast_channel --> tree
```

The first two associations indicate that the abstract ports do not appear in the concrete architecture, resulting in a new concrete signature for the parser and the analyzer. This change in signature reflects the difference between port-to-port communication and shared-memory communication by direct reading and writing of a shared tree. As an analogous example, consider two procedures that communicate through direct calls. If we reimplement this architecture so that the procedures communicate only indirectly through a shared variable, the signature of both procedures would change. The third association says that dataflow connection is implemented by the abstract syntax tree.⁴

We are interested in three specific questions:

- Does the concrete shared-memory architecture implement the abstract dataflow architecture under the completeness assumption and with respect to a given mapping between architectures?
- Is the mapping between the two architectures meaningful? A relative correctness proof is only as meaningful as the mapping between architectures.
- Assuming that the shared-memory implementation of dataflow is correct, under what conditions can it be composed with correct implementations of other parts of the compiler to form a correct and complete compiler architecture?

The running examples in the paper provide a detailed answer to each of these questions.

⁴The tree is a component. A component can be used to implement other components, or it can be used in conjunction with connectors to implement a connection.

4 Formal Criterion of Correctness

Because of the completeness assumption, we must prove not only that a concrete architecture does not lose properties of the abstract architecture, but also that no new properties about the abstract architecture can be inferred from the concrete architecture. There are standard mathematical concepts that can be used for this purpose.

An *interpretation mapping* is an association between the constants, functions, and predicates of an abstract and a concrete theory. An interpretation mapping is called a *theory interpretation* if the mapped axioms of the abstract theory become theorems of the concrete theory. Note that theory interpretation is just Hoare's approach to reasoning about the correctness of implementations [9]. We additionally require that, if a sentence is not in the abstract theory, its image is not in the concrete theory.

Let Θ and Θ' be theories associated with an abstract and a concrete architecture, respectively. Let I be an interpretation mapping from Θ to Θ' . Then, we must have, for every sentence F ,

$$\text{if } F \in \Theta \text{ then } I(F) \in \Theta'$$

for I to be a theory interpretation.

Since we require that an architecture be complete with respect to a given level of detail, we additionally must know that the concrete architecture adds no new facts about the abstract architecture. Therefore, we require that

$$\text{if } F \notin \Theta \text{ then } I(F) \notin \Theta'$$

This says that, if a sentence is not in the abstract theory, its image cannot be in the concrete theory. A theory interpretation I having this property is said to be a *faithful interpretation*. Observe that Θ' is a conservative extension of Θ provided the identity map faithfully interprets Θ in Θ' .

Note that a concrete architecture can contain facts not related to the abstract architecture. Therefore, a concrete architecture can introduce new styles and new objects. For example, a concrete architecture may introduce a specification for part of the runtime environment, such as a wrapper for remote procedure calls that will replace the standard one provided by the operating system.

5 First-Order Architectures

We want to leave open the choice of language for specifying an architecture. Therefore, we represent architectures as first-order theories, but our correctness and composition results in no way depend on this choice.

The representation of the dataflow and the shared-memory architectures in Figures 1 and 3, respectively, depend on the styles used in their construction. The dataflow-style vocabulary contains predicates for describing functional components, ports, values associated with ports, dataflow channels, values associated with dataflow channels, and connections of channels to ports. More precisely, the following sorts denote the first-class objects in a dataflow theory: *channel*, *function*, *iport*, and *oport*. We also make use of sorts *bool* and *val*, where *val* denotes the set of all possible values. The dataflow style has the following operations.

OutPort:	<i>oport</i> × <i>function</i> → <i>bool</i>
Supplies:	<i>oport</i> × <i>val</i> → <i>bool</i>
InPort:	<i>iport</i> × <i>function</i> → <i>bool</i>
Accepts:	<i>iport</i> × <i>val</i> → <i>bool</i>
Carries:	<i>channel</i> × <i>val</i> → <i>bool</i>
Connects:	<i>channel</i> × <i>oport</i> × <i>iport</i> → <i>bool</i>

The number of functions, ports, and channels that can appear in a particular architecture is unbounded. We do not bother to state the general well-formedness axioms associated with this style, or with others. An example of a general dataflow axiom is that every function must have at least one port.

The shared-memory style uses the reading and writing of a variable for intercommunication. Shared-variable communication is modeled using a call site as an interface between a function and the shared variable.⁵ A call site serves the same purpose as a port in the dataflow style. The name of every different call site must be unique. The shared-memory style has the following style-specific sorts: *variable* denotes the set of all possible variables and *site* denotes the set of all possible call sites of which there are two kinds. The sort *rsite* denotes the sites that read, or input, values; the sort *wsite* denotes the ones that write, or output, values. The signature for the shared-memory style is

Holds:	<i>variable</i> × <i>val</i> → <i>bool</i>
CallSite:	<i>site</i> × <i>function</i> → <i>bool</i>
Writes:	<i>wsite</i> × <i>variable</i> → <i>bool</i>
Puts:	<i>wsite</i> × <i>val</i> → <i>bool</i>
Reads:	<i>rsite</i> × <i>variable</i> → <i>bool</i>
Gets:	<i>rsite</i> × <i>val</i> → <i>bool</i>

Table 1 contains (partial) theories associated with the two architectures in Figures 1 and 3. Θ_D denotes the dataflow theory and Θ_M the shared-memory theory. Dataflow theory Θ_D says that the parser and analyzer are functional components, the parser's output port can supply values of type *ast*, the analyzer's input

⁵We could have chosen not to model call sites or some equivalent interface object. We made the decision in order to simplify the style mapping from dataflow to shared-memory.

port can accept values of type *ast*, the dataflow channel can transmit values of type *ast*, and the channel is wired to the ports. The shared-memory theory Θ_M replaces ports with call sites, introduces a variable that can hold values of type *ast*, and employs read and write operations on the variable.

6 Mappings

It is useful to distinguish between two kinds of mappings.

- An *name mapping* associates the objects declared in an abstract architecture with objects declared in a concrete architecture.
- A *style mapping* says how the constructs of the abstract-level style can be implemented in terms of the constructs of the concrete-level style. More specifically, it maps all atomic formulas of the abstract-level theory to formulas of the concrete-level theory.

The two are combined to form an interpretation mapping.

6.1 Name Mapping

We saw a specification of the intended associations between the objects in the two architectures earlier. The only difference in the formal mapping is that we introduce the implicit call sites. Let I_N be name mapping

$$\begin{array}{rcl} oast & \mapsto & site_1 \\ iast & \mapsto & site_2 \\ ast_channel & \mapsto & tree \end{array}$$

which relates the two architectures. The domain of a name mapping can be extended to include all abstract-level terms by mapping variables to themselves.

6.2 Style Mapping

Let I_S denote the style mapping in Figure 4 from the dataflow style to the shared-memory style. The t_i denote terms, which in our examples are restricted to logical constants and variables.⁶ The last association specifies the implementation strategy. It says that any instance of *Connects*(t_1, t_2, t_3) can be implemented by having call site t_2 , corresponding to output port t_2 , be the interface point that provides the values used in the writing of variable t_1 , corresponding to channel t_1 . On the receiving end of a transmission, input port and call site t_3 serve the same function. The other associations say that channels are mapped to variables, that output ports are mapped to calls that supply values, and that

⁶Note that our languages contain no function symbols. A treatment of them can be found in [6].

input ports are mapped to calls that receive values. The *Puts* and *Gets* predicates ensure that the right kind of site is associated with the each kind of port.

6.3 Interpretation Mapping

An *interpretation mapping* I is determined from a name mapping I_N and a style mapping I_S , as follows: for every predicate P , all terms t_1, t_2, \dots, t_n , every variable x , and all formulas F and G of the abstract language,

$$\begin{aligned} I(P(t_1, t_2, \dots, t_n)) &= I_S(P(I_N(t_1), I_N(t_2), \dots, I_N(t_n))) \\ I(\neg F) &= \neg(I(F)) \\ I(F \wedge G) &= I(F) \wedge I(G) \\ I(F \vee G) &= I(F) \vee I(G) \\ I(F \supset G) &= I(F) \supset I(G) \\ I(\forall x F) &= \forall x I(F)^7 \\ I(\exists x F) &= \exists x I(F) \end{aligned}$$

Let I_M^D denote the interpretation mapping from theory Θ_D to theory Θ_M . Both the ground facts and general axioms in Θ_D must be mapped. For example,

$$\begin{aligned} I_M^D(\text{Connects}(ast_channel, oast, iast)) &= I_S(\text{Connects}(I_N(ast_channel), \\ &\quad I_N(oast), I_N(iast))) \\ &= I_S(\text{Connects}(tree, site_1, site_2)) \\ &= \text{Writes}(site_1, tree) \wedge \text{Reads}(site_2, tree) \end{aligned}$$

which is the intended implementation.

7 Proof Obligations

A relative correctness proof involves two steps. First, we must prove the correctness of the relevant style mapping. The proof is performed only once; it need not be repeated when the two styles are used. Second, we must demonstrate the relative correctness of the two architectures with respect to the interpretation mapping formed using the two styles.

7.1 Proof of a Style Mapping

The crucial part of the proof is concerned with the validity of the connector mapping. We would like to know that a dataflow connection can be implemented by the reading and writing of a shared memory location, which is modeled as a variable. This requires a definition of the semantics of both forms of connection. We choose an axiomatic style of semantic definition suitable for describing both safety and fairness properties.

⁷In general, the range of quantifiers must be restricted to a subset of the concrete domain, see [6]. But no restriction is required for our example, because every concrete-level object implements an abstract-level object.

Θ_D	Θ_M
$Function(parser)$	$Function(parser)$
$Function(analyzer)$	$Function(analyzer)$
$OutPort(oast, parser)$	$Variable(tree)$
$\forall v[Supplies(oast, v) \supset ast(v)]$	$\forall v[ast(v) \supset Holds(tree, v)]$
$InPort(iast, analyzer)$	$CallSite(site_1, parser)$
$\forall v[ast(v) \supset Accepts(iast, v)]$	$\forall v[Puts(site_1, v) \supset ast(v)]$
$Channel(ast_channel)$	$Writes(parser, tree)$
$\forall v[ast(v) \supset Carries(ast_channel, v)]$	$CallSite(site_2, analyzer)$
$Connects(ast_channel, oast, iast)$	$\forall v[ast(v) \supset Gets(site_2, v)]$
	$Reads(analyzer, tree)$

Table 1: Partial Dataflow and Shared-Memory Theories

$$\begin{aligned}
Function(t_1) &\mapsto Function(t_1) \\
OutPort(t_1, t_2) &\mapsto CallSite(t_1, t_2) \wedge \exists v Puts(t_1, v) \\
Supplies(t_1, t_2) &\mapsto Puts(t_1, t_2) \\
InPort(t_1, t_2) &\mapsto CallSite(t_1, t_2) \wedge \exists v Gets(t_1, v) \\
Accepts(t_1, t_2) &\mapsto Gets(t_1, t_2) \\
Channel(t_1) &\mapsto Variable(t_1) \\
Carries(t_1, t_2) &\mapsto Holds(t_1, t_2) \\
Connects(t_1, t_2, t_3) &\mapsto Writes(t_2, t_1) \wedge Reads(t_3, t_1)
\end{aligned}$$

Figure 4: A Style Mapping

In particular, we use a temporal logic, called the Temporal Logic of Actions (TLA) [11], to define dataflow and shared-memory communication:

- The semantics of dataflow places minimal restrictions on communication. It says that a multiset of values is transmitted between components. Values can be “lost” and out of order. The fairness condition is that eventually a send or receive occurs unless both are impossible. One reason for impossibility could be failure of the communications line.
- The semantics of shared memory requires that transmission preserve ordering and that values cannot be lost. The fairness condition is that all values written into shared memory will eventually be read from the memory if it is possible to read them.

For comparison purposes, the appendix contains an operational definition of the two forms of communication in standard CSP [8], following Allen and Garlan [2]. CSP can be used to model the safety properties, but not the fairness properties.

We formalize the semantics of dataflow and shared-memory connections as TLA theories. We define an interpretation mapping \mathcal{I}_M^D from the dataflow semantics to the shared-memory semantics and show that it is a theory interpretation. This is sufficient to establish

that dataflow can be implemented with a single shared memory location and that, if the shared-memory communication is fair, the dataflow communication is fair.

We make use of the following TLA notation.

Notation	Meaning
f	list of variables in the old state
f'	list of variables in the new state
\mathcal{A}	action—relation between old and new states
$Enabled$	possible to perform action
$[\mathcal{A}]_f$	$\mathcal{A} \vee (f' = f)$
$\langle \mathcal{A} \rangle_f$	$\mathcal{A} \wedge (f' \neq f)$
$\square F$	always F
$\diamond F$	$\neg \square \neg F$ (sometimes F)
$WF_f(\mathcal{A})$	$\square \diamond \langle \mathcal{A} \rangle_f \vee \square \diamond \neg Enabled \langle \mathcal{A} \rangle_f$

The last line says that eventually action \mathcal{A} must either be taken or become impossible to take. For example, a precondition for execution may not be satisfiable.

In the proof, we make use of two TLA inference rules.

$$\text{STL4. } \frac{F \supset G}{\square F \supset \square G}$$

where F and G are temporal formulas, says that, if F implies G , the always F implies always G .

$$\text{TLA2. } \frac{[\mathcal{A}]_f \supset [\mathcal{B}]_g}{\square [\mathcal{A}]_f \supset \square [\mathcal{B}]_g}$$

is a simplification of Lamport's TLA2 axiom that suffices for our purposes. It says that, if action \mathcal{A} implies \mathcal{B} , then always \mathcal{A} implies always \mathcal{B} .

Figures 5 and 6 contain the TLA theories of dataflow and shared-memory, respectively. The quoted boldface symbols are logical constants. In Figure 5, the dataflow connector is denoted by the *flow* state function, which is a multiset, with three operators: *with* is the insertion operator, *less* is the deletion operator, and *choose* is used to select an element from a nonempty multiset. Values carried by the connector must be in set *Type*, the set of all possible values. The dataflow semantic theory is defined to be Φ , which says three things: the dataflow has to start in the initial state, it must always be possible to perform a send or a receive operation, and the communication line eventually responds to send and receive requests if it is possible to do so (fairness). The shared-memory semantic theory, called Ψ , is defined in a similar manner.

$Init_{\Phi}$	$\stackrel{\text{def}}{=}$	$ev = \text{"ready"}$
		$\wedge flow = \text{"emptybag"}$
S_{sender}	$\stackrel{\text{def}}{=}$	$ev = \text{"ready"}$
		$\wedge ev' = \text{"send"}$
		$\wedge flow' = flow$
		$\wedge val' \in \text{Type}$
R_{receiver}	$\stackrel{\text{def}}{=}$	$ev = \text{"ready"}$
		$\wedge ev' = \text{"receive"}$
		$\wedge flow' = flow$
		$\wedge val' = val$
S_{flow}	$\stackrel{\text{def}}{=}$	$ev = \text{"send"}$
		$\wedge ev' = \text{"ready"}$
		$\wedge flow' = \text{flow with } val'$
		$\wedge val' = val$
R_{flow}	$\stackrel{\text{def}}{=}$	$ev = \text{"receive"}$
		$\wedge ev' = \text{"ready"}$
		$\wedge flow \neq \text{"emptybag"}$
		$\wedge val' = \text{choose}(\text{flow})$
		$\wedge flow' = \text{flow less } val'$
N_{flow}	$\stackrel{\text{def}}{=}$	$S_{\text{flow}} \vee R_{\text{flow}}$
N	$\stackrel{\text{def}}{=}$	$N_{\text{flow}} \vee S_{\text{sender}} \vee R_{\text{receiver}}$
w	$\stackrel{\text{def}}{=}$	$(ev, val, flow)$
Φ	$\stackrel{\text{def}}{=}$	$(\exists \text{flow})(Init_{\Phi} \wedge \square[N]_w \wedge WF_w(N_{\text{flow}}))$

Figure 5: Semantics of Dataflow

Interpretation mapping \mathcal{I}_M^D maps constants, state functions, and operators of the dataflow semantics to those of the shared-memory semantics. \mathcal{I}_M^D is defined by

$Init_{\Psi}$	$\stackrel{\text{def}}{=}$	$op = \text{"ready_write"}$
		$\wedge mem = \text{"undefined"}$
W_{writer}	$\stackrel{\text{def}}{=}$	$op = \text{"ready_write"}$
		$\wedge op' = \text{"write"}$
		$\wedge mem' = mem$
		$\wedge val' \in \text{Type}$
R_{reader}	$\stackrel{\text{def}}{=}$	$op = \text{"ready_read"}$
		$\wedge op' = \text{"read"}$
		$\wedge mem' = mem$
		$\wedge val' = val$
W_{mem}	$\stackrel{\text{def}}{=}$	$op = \text{"write"}$
		$\wedge op' = \text{"ready_read"}$
		$\wedge mem' = val'$
		$\wedge val' = val$
R_{mem}	$\stackrel{\text{def}}{=}$	$op = \text{"read"}$
		$\wedge op' = \text{"ready_write"}$
		$\wedge mem \neq \text{"undefined"}$
		$\wedge mem' = val'$
		$\wedge val' = mem$
M_{mem}	$\stackrel{\text{def}}{=}$	$W_{\text{mem}} \vee R_{\text{mem}}$
M	$\stackrel{\text{def}}{=}$	$M_{\text{mem}} \vee W_{\text{writer}} \vee R_{\text{reader}}$
u	$\stackrel{\text{def}}{=}$	(op, val, mem)
Ψ	$\stackrel{\text{def}}{=}$	$Init_{\Psi} \wedge \square[M]_u \wedge WF_u(M_{\text{mem}})$

Figure 6: Semantics of Shared Memory

ev	\mapsto	op
$flow$	\mapsto	mem
"emptybag"	\mapsto	"undefined"
"ready"	\mapsto	$\text{either}(\text{"ready_write"}, \text{"ready_read"})$
"send"	\mapsto	"write"
"receive"	\mapsto	"read"
$t_1 \text{ with } t_2$	\mapsto	t_2
$t_1 \text{ less } t_2$	\mapsto	t_2
$\text{choose}(t_1)$	\mapsto	t_1

where t_1 and t_2 are terms. The last three associations interpret multiset operations in the context of our specific weak fairness condition on shared memory.

To show that \mathcal{I}_M^D is a theory interpretation, we need to prove that $\Psi \supset \mathcal{I}_M^D(\Phi)$. The first step is to prove that

$$Init_{\Psi} \supset \mathcal{I}_M^D(Init_{\Phi}). \quad (1)$$

Applying \mathcal{I}_M^D to $Init_{\Phi}$ we get:

$$\begin{aligned} op &= \text{either}(\text{"ready_write"}, \text{"ready_read"}) \\ \wedge mem &= \text{"undefined"}. \end{aligned}$$

Hence (1) holds. The second step is to show that

$$\square[M]_u \supset \mathcal{I}_M^D(\square[N]_w). \quad (2)$$

We can easily show that

$$\mathcal{W}_{writer} \supset \mathcal{I}_M^D(\mathcal{S}_{sender}) \quad (3)$$

$$\mathcal{R}_{reader} \supset \mathcal{I}_M^D(\mathcal{R}_{receiver}) \quad (4)$$

$$\mathcal{W}_{mem} \supset \mathcal{I}_M^D(\mathcal{S}_{flow}) \quad (5)$$

$$\mathcal{R}_{mem} \supset \mathcal{I}_M^D(\mathcal{R}_{flow}) \quad (6)$$

from which we infer that

$$[\mathcal{M}]_u \supset \mathcal{I}_M^D([\mathcal{N}]_w).$$

Applying rule TLA2, we conclude that (2) holds. The third step is to show that

$$WF_u(\mathcal{M}_{mem}) \supset \mathcal{I}_M^D(WF_w(\mathcal{N}_{flow})). \quad (7)$$

From (3)–(6), we get

$$(\mathcal{M}_{mem})_u \supset \mathcal{I}_M^D((\mathcal{N}_{flow})_w).$$

Applying rule STL4 twice and the definition of \Diamond , we get

$$\Box\Diamond(\mathcal{M}_{mem})_u \supset \Box\Diamond\mathcal{I}_M^D((\mathcal{N}_{flow})_w).$$

From the definition of *Enabled*, we have

$$Enabled \mathcal{I}_M^D((\mathcal{N}_{flow})_w) \supset Enabled (\mathcal{M}_{mem})_u.$$

Since

$$\mathcal{I}_M^D(Enabled (\mathcal{N}_{flow})_w) \supset Enabled \mathcal{I}_M^D((\mathcal{N}_{flow})_w),$$

we apply rule STL4 to get

$$\Box\Diamond\neg Enabled (\mathcal{M}_{mem})_u \supset \Box\Diamond\neg\mathcal{I}_M^D(Enabled (\mathcal{N}_{flow})_w),$$

from which we conclude that fairness condition (7) holds.

7.2 Relative Correctness Proof

We must show that I_M^P is a theory interpretation and that it is faithful. A proof of the former is straightforward. For example, under I_M^P the axiom

$$Connects(ast_channel, oast, iast)$$

is interpreted as

$$\begin{aligned} & Writes(parser, tree) \\ & \wedge Reads(analyzer, tree) \end{aligned}$$

which is a theorem that follows directly from Θ_M .

To show faithfulness, notice that I_M^P induces a mapping I' from shared-memory structures to dataflow structures as follows. If I_M^P maps atomic dataflow formula $P(\bar{x})$ to shared-memory formula F , then I' assigns to dataflow predicate P the set of shared-memory tuples that satisfy F .

Given a model D of Θ_D , we can construct a model M of Θ_M as follows. The universe of M is the same as D . The assignment to predicates by M is defined as:

$$\begin{aligned} Function &= \{a \in |D| : D \models Function(a)\} \\ Variable &= \{a \in |D| : D \models Channel(a)\} \\ Writes &= \{(a, b) \in |D|^2 : \exists c, d \in |D| \\ &\quad [D \models OutPort(c, a) \wedge \\ &\quad Connects(b, c, d)]\} \\ &\vdots \end{aligned}$$

By a theorem stated in [15] and proved in [16], the fact that induced mapping I' maps M back to D is enough to conclude that I_M^P is faithful.

8 Composing Architectures

A useful form of architecture composition is illustrated in Figure 7. We want to compose two architectures, called “subsystem A” and “subsystem B”, into a single system architecture. We construct a new architecture with components “A” and “B” connected through new interfaces. If two conditions are satisfied, the three architectures can be combined to form a composite system that is correct if the three subsystems are.




Figure 7: Illustration of Subsystem Composition

Let Θ_1 and Θ_2 be theories that represent two abstract architectures. Let Θ'_1 and Θ'_2 be concrete theories intended to implement Θ_1 and Θ_2 , respectively. Two pairs of architecture theories can be composed only in ways that preserve faithfulness. More precisely, if

$$I_1: \Theta_1 \rightarrow \Theta'_1 \text{ and } I_2: \Theta_2 \rightarrow \Theta'_2$$

are faithful interpretations, then we want

$$I_1 \cup I_2: \Theta_1 \cup \Theta_2 \rightarrow \Theta'_1 \cup \Theta'_2$$

to be a faithful interpretation. (The union of two theories is the deductive closure of the set-theoretic union of the theories.)

This property holds provided two general conditions are satisfied.

1. The composite interpretation mapping must be a function. For a sentence F , we require that

$$\forall F \in \Theta_1 \cap \Theta_2 [I_1(F) = I_2(F)]$$

which guarantees that interpretation mappings I_1 and I_2 agree on shared objects and shared style constructs.

2. It must not be possible to infer new facts about the composite abstract architecture from the composite concrete architecture. That is, for language L_1 of Θ_1 and L_2 of Θ_2 , if

$$F \text{ is a sentence of } L_1 \cup L_2$$

and

$$\Theta'_1 \cup \Theta'_2 \vdash I(F)$$

then we must prove that

$$I(\Theta_1) \cup I(\Theta_2) \vdash I(F).$$

The intuition behind the second condition can be illustrated by means of a simple example. Consider an architecture in which there is a dataflow connection from A to B and another architecture that has dataflow connection from B to C . Suppose that both flows are implemented correctly in concrete architectures, but that in one A writes some variable x and in the other C reads a variable x . Each implementation is correct, since neither introduces a new dataflow. However, the composite concrete architecture reads and writes x , from which we can infer an entirely new abstract dataflow connection from A to C . Consequently, the composite abstract architecture is not faithfully interpreted (by the composite mapping) in the composite concrete architecture (under the original assumption that dataflow is intransitive).

Although the second condition is a rather strong logically, it appears to be flexible enough for architecture composition. The form of composition illustrated in Figure 7 can be handled easily by allowing two abstract architectures to share only one component and possibly its interface points. Styles can be shared but no other objects. These constraints guarantee that the two conditions above are satisfied, and the desired composition can be performed in two steps.

Another useful form of composition is the chaining together of a sequence of correct architectures. Since faithful interpretation is transitive, intermediate architectures can be omitted in the development of a concrete architecture. Intermediate architectures arise because we make explicit all important intermediate steps in a development, even if they correspond to small architectural changes. The intermediate architectures need not be explicit as long as there is a sequence of instances

of refinement patterns that connect the first (most abstract) and last (most concrete) architectures in the sequence.

We return to the compiler architecture in Figure 2 to give a specific example of composition. We proved that the dataflow connection between the parser and the analyzer is implemented correctly by means of the reading and writing of the tree. That is, we showed that dataflow theory Θ_D is implemented correctly by theory Θ_M with respect to mapping I_M^P . Similarly, we can show that the dataflow connection from the lexical analyzer to the parser is correctly implemented by the pipeline connection in the concrete architecture. The two abstract-concrete pairs of architectures share a common component, the parser, but no interface points. Therefore, our second condition is satisfied and we can compose the two pairs directly. (The two mappings are constructed to meet the first condition.) No linking architecture is needed.

9 Related Work

The utility of architecture hierarchies was recognized in the 1970s, but architecture hierarchy was studied only informally at that time. Several notations were developed for describing architectures, including those of Jackson [10], Yourdan and Constantine [17], and DeMarco [5], but little attention was given to understanding the relationship between levels of abstraction.

Moriconi and Hare [14] formalized a relationship between levels in a hierarchy and used the technique of Hoare [9] to prove the relative correctness of two stylistically different architectures. Hoare's technique involves a proof of only theory interpretation, and not of faithfulness. They were the first to introduce a completeness assumption for architectures. An architecture was allowed to contain only finitely many objects (constants), which enabled them to fully mechanize correctness proofs. The completeness assumption, as formalized in this paper, applies equally well to infinite architectures. For example, it is possible to quantify over infinite types (such as integers) and to reason about dynamic architectures with an unbounded number of processes.

The technique of Hoare has been applied more recently to architecture by Broy [4], Brinksma [3], and others. Broy's component refinements turn out to be conservative because interface signatures are preserved, but his connection refinements may not be because additional flows could be added to a channel. Brinksma justifies channel splitting on the basis of behavioral reasoning; application of his rule can violate the completeness assumption.

A Hoare-style representation mapping has been applied to dynamic architectures by Luckham et al [12, 13]. A language called Rapide is used to define executable ar-

chitectures based on distributed event processing. Mappings relate concrete events to abstract events and are used as the basis for comparative simulation, a technique that complements ours.

The problem of composition of specifications has been studied in a general semantic framework by Abadi and Lamport [1]. Their results are applicable to any domain, whereas our results are syntactic and specialized to the domain of software architecture. The advantage of a syntactic constraint is that it can be checked easily. The disadvantage is that it is more restrictive than semantic composition. Broy [4] gives three operators for composing functional-style architectures, but does not consider the composition of architectures involving multiple styles.

10 Conclusion

An architecture for a large, complex system, and even some simple systems, will involve multiple levels of detail expressed in multiple architectural styles. The novel contributions of the work reported here are:

- A formal criterion for proving that one architecture implements another architecture, even if they are described in different architectural styles. A change in the representation of a component, an interface, or a connector is handled, but a change in the representation of a type requires a slightly different criterion.
- A decomposition of the mapping between architectures into type-level properties that are proved once for every pair of styles and instance-level properties that are proved for every pair of architectures. The importance of this decomposition was underscored by a proof that the connectors of a common concrete style implement the connectors of a common abstract style. The proof was somewhat complicated, establishing both safety and fairness properties, but it does not need to be repeated each time the styles are used.
- Syntactic criteria for composing architectures such that the composition of two correct architectures is correct. One specific composition operator, which is useful for putting together subsystems, allows two architectures to be composed provided they share only components and their interface points. Another composition operator is used to eliminate intermediate levels in an architecture hierarchy.

Our approach applies to any logic used to represent an architecture; it does not depend on a particular architecture definition language or a particular kind of connector semantics. A more comprehensive treatment of

the formal techniques in this paper can be found in a companion paper [15].

The work reported here may have implications in several subareas of software-architecture research.

- **Language design.** An architecture definition language (ADL) should treat all refineable objects, including components, interface points, and connectors, as first-class in the sense that they should be named objects with independent meaning. Another implication is that an ADL should make it impossible to subvert the completeness assumption. For example, an ADL type system should not allow components to be values, which would allow interactions to be created indirectly. The last implication is that an ADL should support the specification of two kinds of mappings: style mappings and name mappings between architectures.
- **Refinement methodology.** It seems clear that after-the-fact proof of an architecture hierarchy will be very difficult. This is true primarily because of the need to establish conservativeness (modulo renaming). An incremental development strategy that minimizes the number and difficulty of architecture-specific proofs is needed. One candidate approach involving correctness-preserving architectural transformations is described in [15].
- **Style design.** Styles are an important vehicle for organizing reusable architectural design information. We showed that the specification of style mappings is a key element of style design, and that the semantics of a style can be affected by how the style is intended to be used in relation to other styles.

REFERENCES

- [1] M. Abadi and L. Lamport, "Composing Specifications", *ACM Transactions on Programming Languages and Systems*, Vol. 15, No. 1, January 1993, pp. 73–132.
- [2] R. Allen and D. Garlan, "Formalizing Architectural Connection", *Proceedings of the Sixteenth International Conference on Software Engineering*, May 1994, pp. 71–80.
- [3] E. Brinksma, B. Jonsson, and F. Orava, "Refining Interfaces of Communicating Systems", *TAP-SOFT'91: Lecture Notes in Computer Science 494*, S. Abramsky and T.S.E. Maibaum, Eds., Springer-Verlag, 1991, pp. 297–312.
- [4] M. Broy, "Compositional Refinement of Interactive Systems", No. 89, Digital Systems Research Center, Palo Alto, California, July 1992.

- [5] T. DeMarco, *Structured Analysis and System Specification*, Yourdan Press, 1979.
- [6] H. B. Enderton, *A Mathematical Introduction to Logic*, Academic Press, 1972.
- [7] D. Garlan and M. Shaw, "An Introduction to Software Architecture", In *Advances in Software Engineering and Knowledge Engineering*, Volume 1, V. Ambriola and G. Tortora, Eds., World Scientific Publishing Company, 1993.
- [8] C.A.R. Hoare, *Communicating Sequential Processes*, Prentice-Hall, 1985.
- [9] C.A.R. Hoare, "Proof of correctness of data representations", *Acta Informatica*, Vol. 1, No. 4, 1972, pp. 271-281.
- [10] M.A. Jackson, *Principles of Program Design*, Academic Press, 1975.
- [11] L. Lamport, "The Temporal Logic of Actions", Technical Report 79, Digital Systems Research Center, Palo Alto, California, December 1991. (to appear in *ACM Transactions on Programming Languages and Systems*)
- [12] D.C. Luckham, L.M. Augustin, J.J. Kenney, J.S. Vera, D. Bryan, and W. Mann, "Specification and Analysis of System Architecture Using Rapide", to appear in *IEEE Transactions on Software Engineering*.
- [13] D.C. Luckham, J. Vera, D. Bryan, L. Augustin, and F. Belz", "Partial Orderings of Event Sets and Their Application to Prototyping Concurrent, Timed Systems", *Journal of Systems and Software*, Vol. 21, No. 3, June 1993, pp. 253-265.
- [14] M. Moriconi and D.F. Hare, "The PegaSys System: Pictures as Formal Documentation of Large Programs", *ACM Transactions on Programming Languages and Systems*, Vol. 8, No. 4, October 1986, pp. 524-546.
- [15] M. Moriconi, X. Qian, and R. Riemenschneider, "Correct Architecture Refinement", to appear in *IEEE Transactions on Software Engineering*.
- [16] M. Moriconi, X. Qian, and R. Riemenschneider, "A Formal Approach to Correct Refinement of Software Architectures", Technical Report SRI-CSL-94-13, Computer Science Laboratory, SRI International, Menlo Park, California, August 1994.
- [17] E. Yourdan and L.L. Constantine, *Structured Design: Fundamentals of a Discipline of Computer Program and Systems Design*, Prentice-Hall, Inc., Englewood Cliffs, N.J., 1979.

A Proof of Connector Mapping in CSP

We can define the semantics of the dataflow and shared memory styles in CSP [8], following Allen and Garlan [2]. We make use of the following CSP notation.

<u>Notation</u>	<u>Meaning</u>
αP	the alphabet of process P
$P \parallel Q$	P in parallel with Q
$a \rightarrow P$	a then P
$a \rightarrow P \mid b \rightarrow Q$	a then P choice b then Q ($a \neq b$)
$P \backslash C$	P without C (hiding)
$f : A \rightarrow B$	f is a function mapping A to B

We also make use of the count process CT , defined as follows.

$$\begin{aligned} CT_0 &= (up \rightarrow CT_1 \mid around \rightarrow CT_0) \\ CT_{n+1} &= (up \rightarrow CT_{n+2} \mid down \rightarrow CT_n) \end{aligned}$$

The CSP semantics is essentially the same as the TLA semantics. However, a connector is modeled directly in TLA by a state function. It is modeled indirectly in CSP as a process, which essentially computes the state function. Standard CSP cannot be used to express fairness of the kind in our example. Therefore, we prove only safety.

The CSP semantics for the dataflow style is	
αDFS	= Sender \parallel Receiver \parallel Flow
$\alpha Sender$	= {oport}
Sender	= oport \rightarrow Sender
$\alpha Receiver$	= {iport}
Receiver	= iport \rightarrow Receiver
$\alpha Flow$	= {oport, iport}
Flow	= $(CT_0 \parallel Flow') \setminus \{around, down, up\}$
$\alpha Flow'$	= {around, down, up, oport, iport}
Flow'	= oport \rightarrow up \rightarrow Flow' around \rightarrow Flow' down \rightarrow iport \rightarrow Flow'

and the CSP semantics for the shared-memory style is

αSMS	= Writer \parallel Reader \parallel Var
$\alpha Writer$	= {write}
Writer	= write \rightarrow Writer
$\alpha Reader$	= {read}
Reader	= read \rightarrow Reader
αVar	= {write, read}
Var	= write \rightarrow read \rightarrow Var

We must show that the shared-memory style is a correct implementation of the dataflow style. Intuitively, every behavior of the shared-memory style should correspond to an allowable behavior of the dataflow style. Since the alphabets of the two styles are different, this can be done using the CSP change-of-symbol operator f : $f(write) = oport$ and $f(read) = iport$. Hence, the correctness proof amounts to showing that $f(SMS) \sqsubseteq DFS$, which is straightforward.

DISTRIBUTION LIST

addresses	number of copies
JOSEPH A. CAROZZONI RL/C3CA 525 BROOKS ROAD ROME NY 13441-4505	1
SRI INTERNATIONAL 333 RAVENSWOOD AVENUE MENLO PARK, CA 94025	1
ROME LABORATORY/SUL TECHNICAL LIBRARY 26 ELECTRONIC PKY ROME NY 13441-4514	1
ATTENTION: DTIC-DCC DEFENSE TECHNICAL INFO CENTER 8725 JOHN J. KINGMAN ROAD, STE 0944 FT. BELVOIR, VA 22060-6218	2
ADVANCED RESEARCH PROJECTS AGENCY 3701 NORTH FAIRFAX DRIVE ARLINGTON VA 22203-1714	1
AFIT ACADEMIC LIBRARY/LDCE 2950 P STREET AREA B, BLDG 642 WRIGHT-PATTERSON AFB OH 45433-7765	1
PHILLIPS LABORATORY PL/TL (LIBRARY) 5 WRIGHT STREET HANSOM AFB MA 01731-3004	1
DL AL HSC/HRG, BLDG 190 2698 G STREET WPAFB OH 45433-7604	1